Design-theoretic encoding of deterministic hypotheses as constraints and correlations into U-relational databases

Bernardo Gonçalves
LNCC – National Laboratory
for Scientific Computing
Petrópolis, Brazil
bgonc@lncc.br

Fabio Porto
LNCC – National Laboratory
for Scientific Computing
Petrópolis, Brazil
fporto@lncc.br

arXiv:1411.5196v1 [cs.DB] 19 Nov 2014

ABSTRACT

In view of the paradigm shift that makes science ever more data-driven, in this paper we consider deterministic scientific hypotheses as uncertain data. In the form of mathematical equations, hypotheses symmetrically relate aspects of the studied phenomena. For computing predictions, however, deterministic hypotheses are used asymmetrically as functions. We refer to Simon's notion of structural equations in order to extract the (so-called) causal ordering embedded in a hypothesis. Then we encode it into a set of functional dependencies (fd's) that is basic input to a design-theoretic method for the synthesis of U-relational databases (DB's).

The causal ordering captured from a formally-specified system of mathematical equations into fd's determines not only the constraints (structure), but also the correlations (uncertainty chaining) hidden in the hypothesis predictive data. We show how to process it effectively through original algorithms for encoding and reasoning on the given hypotheses as constraints and correlations into U-relational DB's. The method is applicable to both quantitative and qualitative hypotheses and has underwent initial tests in a realistic use case from computational science.

Categories and Subject Descriptors

H.2.1 [**Information Systems**]: Logical Design

Keywords

Deterministic hypotheses, design by synthesis, U-relations

1. INTRODUCTION

As part of the paradigm shift that makes science ever more data-driven, deterministic scientific hypotheses can be seen as: principles or ideas, which are mathematically expressed and then implemented in a program that is run to give their *decisive* form of data. For a description of the research vision of hypothesis management and its significance, we refer the reader to [10].[1] In this paper we engage in a theoretical exploration on deterministic hypotheses as a kind of uncertain data.

Target applications. Our framework is geared at hypothesis management applications. Examples of structured deterministic hypotheses include tentative mathematical models in physics, engineering and economical sciences, or conjectured boolean networks in biology and social sciences. These are important reasoning devices, as they are solved to generate predictive data for decision making in both science and business. But the complexity and scale of modern scientific problems require proper data management tools for the predicted data to be analyzed more effectively.

Probabilistic DBs. Probabilistic databases (p-DBs) qualify as such tool, as they have evolved into mature technology in the last decade [26]. One of the state-of-the-art probabilsitic data models is the U-relational representation system with its probabilistic world-set algebra (p-WSA) implemented in MayBMS [17]. That is an elegant extension of the relational model we refer to in this paper for the management of uncertain and probabilistic data. Our goal is to develop means to extract a hypothesis specification and encode it into a U-relational DB seamlessly, ensuring consistency and quality w.r.t. the given hypothesis structure. In short, we shall flatten deterministic models into U-relations.

Structural equations. Given a system of equations with a set of variables appearing in them, in a seminal article Simon introduced an asymmetrical, functional relation among variables that establishes a (so-called) *causal ordering* [24]. Along these lines, we shall extract the causal ordering of a deterministic hypothesis and encode it into a set of fd's that is basic input to our synthesis of U-relational DBs. As we shall see, the causal ordering we capture in fd's determines not only the constraints (structure), but also the correlations (uncertainty chaining) hidden in predictive data. In comparison with research on causality in DBs [20] (cf. §6.1), our framework comprises a technique for encoding and processing causality at schema level.

Background theory. We rely on the following body of background theoretical work: (i) U-relations and p-WSA [17], as our design-theoretic method is shaped for U-relational DBs; (ii) classical theory of fd's and normalization [1, 27], and Bernstein's design-by-synthesis

[1]Also, to anticipate §6.1, it can be understood in comparison as a specific, shorter-term research path along the lines of Haas et al.'s *models-and-data* program [12].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$15.00.

approach [4]; and finally, (iii) Simon's notion of structural equation models (SEMs) [22, 24].

List of contributions. Overall, this paper presents specific technical developments over the ϒ-DB vision [10]. In short, it shows how to encode deterministic hypotheses as uncertain data, viz., as constraints and correlations into U-relational DBs. Our detailed technical contributions are:

- We study the relationship between SEMs and fd's and present an *encoding scheme* that extracts the causal ordering of a given (formally specified) deterministic hypothesis and encodes it into a set Σ of fd's; we then uncover some of its main properties. In short, given a hypothesis, we show how to transform it algorithmically into a set Σ of fd's in order to design a relational DB over it.
- For a "good" design, we show that the hypothesis causal ordering mapped into fd set Σ needs to be further processed in terms of *acyclic reasoning on its reflexive pseudo-transitive closure*. We present an original, efficient algorithm for that, which returns an fd set Σ' we motivate and define to be the *folding* $\Sigma^{\looparrowright}$ of Σ. Then we apply a variant of Bernstein's synthesis algorithm (say, '4C') to render, given $\Sigma^{\looparrowright}$, a relational schema $\boldsymbol{H}_k = \bigcup_{i=1}^{n} H_k^i$ shown to bear desirable properties for hypothesis management — yet up to the capabilities of a traditional relational DB at this stage of the design pipeline. In short, this is a design-theoretic technique that uses the extracted fd's as constraints.
- Once schema $\boldsymbol{H}_k$ is synthesized, datasets computed from the hypothesis under alternative trials (input settings) can be loaded into it. Finally, then, through a different manipulation on the primitive fd set Σ, we extract the *uncertainty chaining* (correlations) from it into a Σ'' which is, together with $\boldsymbol{H}_k$, input to an original synthesis procedure (say, '4U') to render U-relational $\boldsymbol{Y}_k = \bigcup_{j=1}^{m} Y_k^j$. In short, this is a principled technique to *introduce uncertainty* in p-WSA given a set of fd's.

After introducing notation and basic concepts in §2, we present through §3–§5 the contributions (resp.) listed above. We discuss related work and the applicability of our framework in §6, and also point to a use case scenario from which we have extracted and encoded real-world hypotheses and conducted some initial experiments. Finally, §7 concludes the paper.

2. PRELIMINARIES

As notational conventions, we write X, Y, Z to denote sets of relational attributes and A, B, C to denote single attributes. Also, we write XY as shorthand for $X \cup Y$, and $R[XZ]$ to denote relation R has scheme $U = XZ$ with designated key constraint $X \to Z$.

2.1 U-Relations and Probabilistic WSA

A U-relational DB or U-DB is a finite set of structures,

$$\boldsymbol{W} = \{\langle R_1^1, \ldots, R_m^1, p^{[1]}\rangle, \ldots, \langle R_1^n, \ldots, R_m^n, p^{[n]}\rangle\},$$

of relations $R_1^i, \ldots, R_m^i$ and numbers $0 < p^{[i]} \leq 1$ such that $\sum_{1 \leq i \leq n} p^{[i]} = 1$. An element $R_1^i, \ldots, R_m^i, p^{[i]} \in \boldsymbol{W}$ is a *possible world*, with $p^{[i]}$ being its probability [17].

Probabilistic world-set algebra (p-WSA) consists of the operations of relational algebra, an operation for computing tuple confidence conf, and the repair-key operation for introducing uncertainty — by giving rise to alternative worlds as maximal-subset repairs of an argument key (cf. Def. 1) [17].

Def. 1. *Let $R_\ell[U]$ be a relation, and $XA \subseteq U$. For each possible world $\langle R_1, \ldots, R_m, p\rangle \in \boldsymbol{W}$, let $A \in U$ contain only numerical values greater than zero and let R_ℓ satisfy the fd $(U \setminus A) \to U$. Then, repair-key is:*

$$[\![\textsf{repair-key}_{X@A}(R_\ell)]\!](\boldsymbol{W}) := \left\{\langle R_1, .., R_\ell, R_m, \hat{R}_\ell[U \setminus A], \hat{p}\rangle\right\},$$

where $\langle R_1, ..., R_\ell, R_m, p\rangle \in \boldsymbol{W}$, $\hat{R}_\ell$ is a maximal repair of fd $X \to U$ in R_ℓ, and $\hat{p} = p \cdot \prod_{t \in \hat{R}_\ell} \frac{t.B}{\sum_{s \in R_\ell : s.X = t.X} s.B}$.

U-relations (cf. Fig. 1) have in their schema a set of pairs (V_i, D_i) of *condition columns* (cf. [17]) to map each discrete random variable x_i to one of its possible values (e.g., $\mathsf{x}_1 \mapsto 1$). The world table W stores their marginal probabilities (cf. the notion of *pc-tables* [26, Ch. 2]). For an illustration of the data transformation from certain to uncertain relations, consider query (1) in p-WSA's extension of relational algebra, whose result set is materialized into U-relation Y_0 as shown in (Fig. 1).

$$Y_0 := \pi_{\phi,\upsilon}(\textsf{repair-key}_{\phi@\textsf{Conf}}(H_0)\,). \tag{1}$$

Also, let $R[\overline{V_i\ D_i}\ \mid sch(R)]$, $S[\overline{V_j\ D_j}\ \mid sch(S)]$ be two U-relations, where $R.\overline{V_i\ D_i}$ is the union of all pairs of condition columns $V_i\ D_i$ in R, then operations $[\![\,\sigma_\psi(R)\,]\!]$, $[\![\,\pi_Z(R)\,]\!]$ and $[\![R \times S]\!]$ issued in relational algebra are rewritten in positive relational algebra on U-relations:

$[\![\sigma_\psi(R)]\!] := \sigma_\psi(R[\overline{V_i\ D_i}\ \mid sch(R)])$;

$[\![\,\pi_Z(R)\,]\!] := \pi_{\overline{V_i\ D_i}\ Z}(R)$;

$[\![R \times S]\!] := \pi_{(R.\overline{V_i\ D_i}\ \cup\ S.\overline{V_i\ D_i}) \to \overline{V\ D}\ \cup\ sch(R) \cup sch(S)}(R \bowtie_{R.\overline{V_i\ D_i}\ \text{is consistent with}\ S.\overline{V_j\ D_j}} S)$.

If R and S have k and ℓ pairs of condition columns each, then $[\![R \times S]\!]$ returns a U-relation with $k + \ell$ such pairs. If $k = 0$ or $\ell = 0$ (or both), then R or S (or both) are classical relations, but the rewrite rules above apply accordingly. All that rewriting is parsimonious translation (sic. [17]): the number of algebraic operations does not increase and each of the operations selection, projection and product/join remains of the same kind. Query plans are hardly more complicated than the input queries. In fact, off-the-shelf relational database query optimizers do well in practice.

For a comprehensive overview of U-relations and p-WSA we refer the reader to [17]. In this paper we look at U-relations from the point of view of p-DB design, for which no methodology has yet been proposed. We are concerned in particular with hypothesis management applications [10].

2.2 Design by Synthesis and Normalization

The problem of design by synthesis has long been introduced by Bernstein in purely symbolic terms as follows [4]: given a set U of attribute symbols and a set Σ of mappings of sets of symbols into symbols (the fd's), find

H_0	ϕ	υ	Conf
	1	1	2
	1	2	2
	1	3	1

Y_0	$V \mapsto D$	ϕ	υ
	$x_0 \mapsto 1$	1	1
	$x_0 \mapsto 2$	1	2
	$x_0 \mapsto 3$	1	3

W	$V \mapsto D$	Pr
	$x_0 \mapsto 1$	.4
	$x_0 \mapsto 2$	.4
	$x_0 \mapsto 3$	.2

Figure 1: U-relation generated by the repair-key operation.

a collection $\boldsymbol{R}=\{R_1, R_2, \ldots, R_n\}$ (the relations) of subsets of U and, for each R_i, a subset of R_i (its designated key) satisfying properties: (P1) each $R_i \in \boldsymbol{R}$ is in 3NF; (P2) $\boldsymbol{R}$ completely characterizes Σ; and (P3) the cardinality $|\boldsymbol{R}|$ is minimal.

More generally, the problem of schema design given dependencies considers the following criteria [1, Ch. 11]:

P1′. ($\simeq$ P1). Desirable properties by normal forms;

P2′. ($\simeq$P2). Preservation of dependencies ("meta-data");

P3. The cardinality $|\boldsymbol{R}|$ is minimal (minimize joins);

P4. Preservation of data (the lossless join property).

There is a trade-off between P1′ and P2′, since normal forms that ensure less redundant schemes may lose the property of dependency preservation [1]. In fact, P2′ is important to prevent the DB from the so-called update anomalies, as the fd's in Σ are viewed as integrity constraints to their associated relations [27, p. 398]. Hypothesis management applications [10], however, are OLAP-like and have an ETL-pipeline characterized by batch-, incremental-only updates and large data volumes. Thus we shall trade P2′ for P1′, to favor succintness (as less redundancy as possible) over dependency preservation (recall BCNF, Def. 2). Also, we shall favor P4 as less joins means faster access to data.

Recall from Ullman [27] that an attribute $A \in U$ is said to be *prime* in relation schema $R[U]$ if it is part of some key for $R[U]$. Def. 2 presents the Boyce-Codd normal form (BCNF) and the third normal form (3NF).

DEF. 2. *Let $R[U]$ be a relation scheme over set U of attributes, and Σ a set of fd's on U. We say that:*

*(a) R is in **BCNF** if, for all $\langle X, A\rangle \in \Sigma^+$ with $A \not\subseteq X$ and $XA \subseteq U$, we have $X \to U$ (i.e., X is a superkey for R);*

*(b) R is in **3NF** if, for all $\langle X, A\rangle \in \Sigma^+$ with $A \not\subseteq X$ and $XA \subseteq U$, we have $X \to U$ or A is prime;*

(c) A schema $\boldsymbol{R}$ is in BCNF (3NF) w.r.t. Σ if all of its schemes $R_1, ..., R_n \in \boldsymbol{R}$ are in BCNF (3NF).

We also recall from [27] the lossless join property (Def. 3) and the notion of dependency preservation (Def. 4).

DEF. 3. *Let $R[U]$ be a relational schema synthesized into collection $\boldsymbol{R} = \bigcup_{i=1}^{n} R_i$ and let Σ be an fd set on attributes U. We say that $\boldsymbol{R}$ has a **lossless join** w.r.t. Σ if for every instance r of $R[U]$ satisfying Σ, we have $r = \bowtie_{i=1}^{n} \pi_{R_i}(r)$.*

DEF. 4. *Let Σ be a set of fd's and $\boldsymbol{R}=\{R_1, R_2, ..., R_n\}$ be a relational schema. We say that $\boldsymbol{R}$ **preserves** Σ if the union of all fd's in $\boldsymbol{R}=\{R_1, R_2, \ldots, R_n\}$ implies Σ.*

Design theory and normalization relies on Armstrong's inference rules (or axioms) of (R0) reflexivity, (R1) augmentation and (R2) transitivity, which forms a sound and complete inference system for reasoning over fd's [27]. From R0-R2 one can derive additional rules, viz., (R3) decomposition, (R4) union and (R5) pseudo-transitivity.

R0. If $Y \subseteq X$, then $X \to Y$;
R1. If $X \to Y$, then $XZ \to YZ$;
R2. If $X \to Y$ and $Y \to W$, then $X \to W$;
R3. If $X \to YZ$, then $X \to Y$ and $X \to Z$;
R4. If $X \to Y$ and $X \to Z$, then $X \to YZ$;
R5. If $X \to Y$ and $YZ \to W$, then $XZ \to W$.

Given an fd set Σ, one can obtain Σ^+, the closure of Σ, by a finite application of rules R0-R5. We are concerned with reasoning over an fd set in order to process its 'embedded' causal ordering. The latter, as we shall see in §4, can be performed in terms of reflexive (pseudo-)transitive reasoning. Note that R2 is a particular case of R5 when $Z = \varnothing$, then we shall refer to $\{\text{R0}, \text{R5}\}$ reasoning and understand R2 included. Def. 5 opens up a way to compute Σ^+ efficiently.

DEF. 5. *Let Σ be an fd set on attributes U, with $X \subseteq U$. Then X^+, the **attribute closure** of X w.r.t. Σ, is the set of attributes A such that $\langle X, A\rangle \in \Sigma^+$.*

Bernstein has long given algorithm XClosure (cf. Alg. 7 in §A) to compute X^+ in time that is polynomial in $|\Sigma| \cdot |U|$ [4]. Finally, we shall also make use of the concept of 'canonical' fd sets (also called 'minimal' [27, p. 390]), see Def. 6.

DEF. 6. *Let Σ be an fd set. We say that Σ is **canonical** if:*

(a) each fd in Σ has the form $X \to A$, where $|A| = 1$;

(b) For no $\langle X, A\rangle \in \Sigma$ we have $(\Sigma - \{\langle X, A\rangle\})^+ = \Sigma^+$;

(c) for each fd $X \to A$ in Σ, there is no $Y \subset X$ such that $(\Sigma \setminus \{X \to A\} \cup \{Y \to A\})^+ = \Sigma^+$.

For an fd set satisfying such properties (Def. 6) individually, we say that it is (a) *singleton-rhs*, (b) *non-redundant* and (c) *left-reduced*. It is said to have an attribute A in X that is 'extraneous' w.r.t. Σ if it is *not* left-reduced (Def. 6-c) [19, p. 74]. Finally, an fd $X \to Y$ in Σ is said *trivial* if $Y \subseteq X$.

2.3 SEMs and Causal Ordering

Given a system of mathematical equations involving a set of variables, to build a *structural equation model* (SEM) is, essentially, to establish a one-to-one mapping between equations and variables [24]. That enables further detecting the hidden asymmetry between variables, i.e., their causal ordering.

DEF. 7. *A **structure** is a pair $\mathcal{S}(\mathcal{E}, \mathcal{V})$, where $\mathcal{E}$ is a set of equations over set $\mathcal{V}$ of variables, $|\mathcal{E}| \leq |\mathcal{V}|$, such that:*

(a) In any subset of k equations of the structure, at least k different variables appear;

(b) In any subset of k equations in which r variables appear, $k \leq r$, if the values of any $(r-k)$ variables are chosen arbitrarily, then the values of the remaining k variables can be determined uniquely — finding these unique values is a matter of solving the equations.

DEF. 8. *Let $\mathcal{S}(\mathcal{E},\mathcal{V})$ be a structure. We say that $\mathcal{S}$ is self-contained or* ***complete*** *if $|\mathcal{E}| = |\mathcal{V}|$.*

Complete structures can be solved for unique sets of values of their variables. In this work, however, we are not concerned with solving sets of mathematical equations at all, but with extracting their causal ordering in view of U-relational DB design. Simon's concept of causal ordering has its roots in econometrics studies (cf. [24]) and to some extent has been taken further in AI with a flavor of Graphical Models (GMs) [8, 22, 7]. In this paper we translate the problem of causal ordering into the language of data dependencies, viz., into fd's.

DEF. 9. *Let $\mathcal{S}$ be a structure. We say that $\mathcal{S}$ is* ***minimal*** *if it is complete and there is no complete structure $\mathcal{S}' \subset \mathcal{S}$.*

DEF. 10. *The* ***structure matrix*** *A_S of a structure $\mathcal{S}(\mathcal{E},\mathcal{V})$, with $f_1, f_2, \ldots, f_n \in \mathcal{E}$ and $x_1, x_2, \ldots, x_m \in \mathcal{V}$, is a $n \times m$ matrix of 1's and 0's in which entry a_{ij} is non-zero if variable x_j appears in equation f_i, and zero otherwise.*

Elementary row operations (e.g., row multiplication by a constant) on the structure matrix may hinder the structure's causal ordering and then are not valid in general [24]. This also emphasizes that the problem of causal ordering is not about solving the system of mathematical equations of a structure, but identifying its hidden asymmetries.

DEF. 11. *Let $\mathcal{S}(\mathcal{E},\mathcal{V})$ be a complete structure. Then a* ***total causal mapping*** *over $\mathcal{S}$ is a bijection $\varphi_t : \mathcal{E} \to \mathcal{V}$.*

Simon has informally described an algorithm (cf. [24]) that, given a complete structure $\mathcal{S}(\mathcal{E},\mathcal{V})$, can be used to compute a *partial* causal mapping $\varphi_p : \mathcal{E} \to \mathcal{V}$ from the set of equations to the set of variables. As shown by Dash and Druzdzel [7], the causal mapping returned by Simon's (so-called) Causal Ordering Algorithm (COA) is not *total* when $\mathcal{S}$ has variables that are *strongly coupled*, i.e., can only be determined simultaneously. They also have shown that any total mapping φ_t over $\mathcal{S}$ must be consistent with COA's partial mapping φ_p [7]. The latter is made partial by design (merge strongly coupled variables) to force its induced causal graph G_{φ_p} to be acyclic.

Dash and Druzdzel's work (cf. [7]) is focused on the correcteness of COA, from a GM point of view. Instead, we shall elaborate on COA in purely symbolic terms, towards encoding structures into fd sets and reasoning over them using Armstrong's rewrite rules R0, R5. For extracting a structure's causal ordering into an fd set, we are only concerned with total causal mappings and then shall have to deal with the issue of cyclic fd's in the causal ordering. We shall map (injectively) variables to relational attributes and (bijectively) equations to fd's.

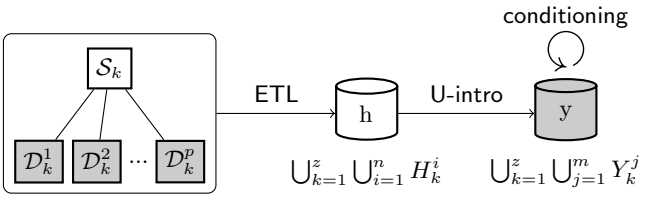

Figure 2: Design-theoretic pipeline for hypothesis encoding.

2.4 Problem Statement

Now we can formulate more precisely the problems in our design pipeline (Fig. 2). In the 'local' view for a hypothesis k, it synthesizes U-relations $\bigcup_{j=1}^{m} Y_k^j$ given its complete structure $\mathcal{S}_k$ and its alternative trial datasets $\bigcup_{\ell=1}^{p} \mathcal{D}_k^\ell$. In fact, the U-intro procedure is operated by the pipeline in the 'global' view of all available hypotheses $k = 1..z$. Their conditioning in the presence of evidence (cf. [10]) is not covered in this paper.

PROBLEM 1. *(**Hypothesis encoding**). Given the (complete) structure $\mathcal{S}_k$ of deterministic hypothesis, extract a total causal mapping φ_t over $\mathcal{S}_k$ and encode φ_t into an fd set Σ_k.*

Following the encoding of a hypothesis structure S into a set Σ of fd's, we target at rendering its relational schema for certainty ('4C,' for short). In short, it is meant to be the minimal-cardinality schema in BCNF that may have a lossless join.

PROBLEM 2. *(**Synthesis '4C'**). Given fd set Σ, derive an fd set Σ' (causal ordering processing) to synthesize a relational schema $\bigcup_{i=1}^{n} H_k^i$ over it satisfying P1′ (BCNF), P3 and striving for P4 while giving up P2′.*

Note in Fig. 2 that coping with these two problems enables the loading of datasets $\bigcup_{\ell=1}^{p} \mathcal{D}_k^\ell$ into schemes $\bigcup_{i=1}^{n} H_k^i$ to accomplish the ETL phase of the design pipeline. The user can then benefit from hypothesis management up to the capabilities of a traditional relational DB. For a full-fledged tool, we shall leverage (globally) the *certain* relations $\bigcup_{k=1}^{z}\bigcup_{i=1}^{n} H_k^i$ to *uncertain* relations $\bigcup_{k=1}^{z}\bigcup_{j=1}^{m} Y_k^j$.

PROBLEM 3. *(**Synthesis '4U'**). Given a collection of relations $\bigcup_{i=1}^{n} H_k^i$ loaded with trial datasets $\bigcup_{\ell=1}^{p} \mathcal{D}_k^\ell$ for each hypothesis k, introduce properly all the uncertainty present in $\bigcup_{k=1}^{z}\bigcup_{i=1}^{n} H_k^i$ w.r.t. encoded fd sets Σ_k for $k=1..z$ into U-relations $\bigcup_{k=1}^{z}\bigcup_{j=1}^{m} Y_k^j$.*

We address problems P1-P3 in the sequel through §3–§5.

3. HYPOTHESIS ENCODING

In this section we present a technique to address Problem 1. For the encoding we shall consider a set Z of attribute symbols such that $Z \simeq \mathcal{V}$, where $\mathcal{S}(\mathcal{E},\mathcal{V})$ is a complete structure; and two special attribute symbols, ϕ, $\upsilon \notin Z$, which are kept to identify (resp.) phenomena and hypotheses. We are explicitly distinguishing symbols in Z, assigned by the user into structure $\mathcal{S}$, from epistemological symbols ϕ and υ. Now, we consider a sense of Simon's into the nature of scientific modeling and interventions [24], summarized in Def. 12.

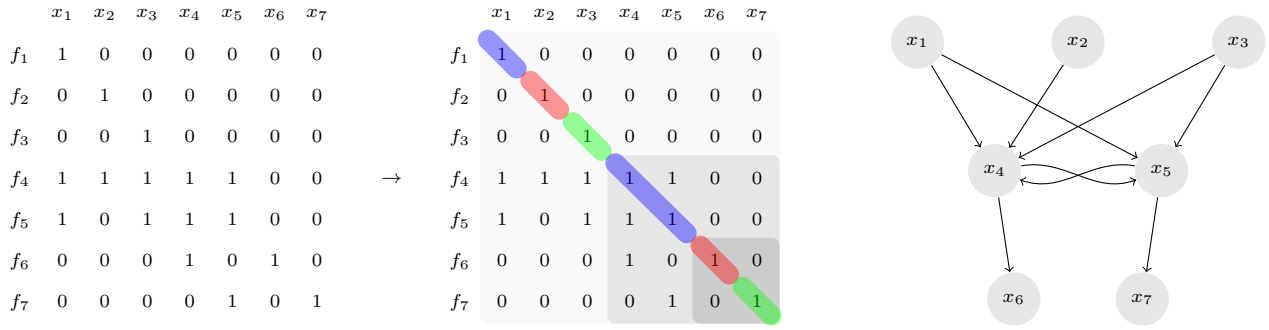

(a) Structure matrix as given. (b) COA_t execution in 3 recursive steps. (c) Directed causal graph G_{φ_t}.

Figure 3: Running Simon's Causal Ordering Algorithm (COA) on a given structure matrix (Fig. 3a). Minimal subsets detected in each recursive step (highlighted in different shades of gray) have their diagonal elements colored (Fig. 3b).

DEF. 12. *Let $\mathcal{S}(\mathcal{E},\mathcal{V})$ be a structure and $x_\ell \in \mathcal{V}$ be a variable. We say that x_ℓ is **exogenous** if there exists an equation $f_k \in \mathcal{E}$ that can be written $f_k(x_\ell) = 0$, i.e., $A_S(k,j) = 1$ iff $j = \ell$. We say that x_ℓ is **endogenous** otherwise.*

Remark 1 introduces an interpretation of Def. 12 with a data dependency flavor.

REMARK 1. *The value of exogenous variables (attributes) is determined empirically, outside of the system (proposed structure $\mathcal{S}$). Such values are, therefore, dependent on the phenomenon id ϕ only. The value of endogenous variables (attributes) is in turn determined theoretically, within the system. They are dependent on the hypothesis id υ and shall be dependent on the phenomenon id ϕ as well.* □

We give (Alg. 1) COA_t, which is a (more detailed) variant of Simon (and Dash-Druzdzel)'s COA. It returns a total causal mapping φ_t, instead of a partial causal mapping. We illustrate it through Example 1 and Fig. 3.

Algorithm 1 COA_t as a variant of Simon's COA.

procedure $\mathsf{COA}_t(\mathcal{S}$: structure over $\mathcal{E}$ and $\mathcal{V})$
Require: $\mathcal{S}$ given is complete, i.e., $|\mathcal{E}| = |\mathcal{V}|$
Ensure: Returns total causal mapping $\varphi_t : \mathcal{E} \to \mathcal{V}$
  $\varphi_t \leftarrow \varnothing$, $\mathcal{S}_c \leftarrow \varnothing$
  for all minimal $\mathcal{S}' \subset \mathcal{S}$ **do**
    $\mathcal{S}_c \leftarrow \mathcal{S}_c \cup \mathcal{S}'$ ▷ store minimal structures in $\mathcal{S}$
    $\mathcal{V}' \leftarrow \mathcal{S}'(\mathcal{V})$
    for all $f \in \mathcal{S}'(\mathcal{E})$ **do**
      $x \leftarrow$ any $x_a \in \mathcal{V}'$
      $\varphi_t \leftarrow \varphi_t \cup \langle f, x\rangle$
      $\mathcal{V}' \leftarrow \mathcal{V}' \setminus \{x\}$
  $\mathcal{T} \leftarrow \mathcal{S} \setminus \bigcup_{\mathcal{S}' \in \mathcal{S}_c} \mathcal{S}'$
  if $\mathcal{T} \neq \varnothing$ **then**
    return $\varphi_t \cup \mathsf{COA}_t(\mathcal{T})$
  return φ_t

EXAMPLE 1. *Consider structure $\mathcal{S}(\mathcal{E},\mathcal{V})$ whose matrix is shown in Fig. 3a. Note that $\mathcal{S}$ is complete, since $|\mathcal{E}| = |\mathcal{V}| = 7$, but not minimal. The set of all minimal subsets $\mathcal{S}' \subset \mathcal{S}$ is $\mathcal{S}_c = \{\{f_1\}, \{f_2\}, \{f_3\}\}$. By eliminating the variables identified at recursive step k, a smaller structure $\mathcal{T} \subset \mathcal{S}$ is derived. Compare the partial causal mapping eventually returned by COA, $\varphi_p \supset \{\langle\{f_4, f_5\}, \{x_4, x_5\}\rangle\}$, to the total causal mapping returned by COA_t, $\varphi_t \supset \{\langle f_4, x_4\rangle, \langle f_5, x_5\rangle\}$. Since x_4 and x_5 are strongly coupled (see Fig.3b), COA_t maps them arbitrarily (i.e., it could be $f_4 \mapsto x_5$, $f_5 \mapsto x_4$ instead). Such total mapping φ_t renders a cycle in the directed causal graph G_{φ_t} (see Fig.3c).* □

We encode complete structures into fd sets by means of (Alg. 2) h-encode. Fig. 4 (left) presents an fd set defined $\Sigma \triangleq \mathsf{h\text{-}encode}(\mathcal{S})$, where $\mathcal{S}$ is shown in Fig. 3. Next we study the main properties of the encoded fd sets.

Algorithm 2 Hypothesis encoding.

procedure H-ENCODE($\mathcal{S}$: structure over $\mathcal{E}$ and $\mathcal{V}$)
Require: $\mathcal{S}$ given is a complete structure, i.e., $|\mathcal{E}| = |\mathcal{V}|$
Ensure: Returns a non-redundant fd set Σ
  $\Sigma \leftarrow \varnothing$
  $\varphi_t \leftarrow \mathsf{COA}_t(\mathcal{S})$
  for all $\langle f_k, x_\ell\rangle \in \varphi_t$ **do**
    $Z \leftarrow x_j$ for all j such that $A_S(k,j) = 1$
    if $|Z| = 1$ **then** ▷ x_ℓ is exogenous
      $\Sigma \leftarrow \Sigma \cup \langle\{\phi\}, \{x_\ell\}\rangle$
    else ▷ x_ℓ is endogenous
      $\Sigma \leftarrow \Sigma \cup \langle Z \setminus \{x_\ell\} \cup \{\upsilon\},\ \{x_\ell\}\rangle$
  return Σ

LEMMA 1. *Let Σ be a singleton-rhs fd set on attributes U. Then $\langle X, A\rangle \in \Sigma^+$ with $XA \subseteq U$ only if $A \subseteq X$ or there is non-trivial $\langle Y, A\rangle \in \Sigma$ for some $Y \subset U$.*

PROOF. See Appendix, §B.1. □

THEOREM 1. *Let Σ be an fd set defined $\Sigma \triangleq \mathsf{h\text{-}encode}(\mathcal{S})$ for some complete structure $\mathcal{S}$. Then it is non-redundant but may not be canonical.*

PROOF. We show that properties (a-b) of Def. 6 must hold for Σ produced by (Alg. 2) h-encode, but property (c) may not hold (i.e., encoded fd set Σ may not be left-reduced). See Appendix, §B.2. □

Finally, it shall be convenient to come with a notion of *parsimonious* fd sets (see Def. 13). This is bit stronger an assumption than canonical fd sets, yet provably just fit to our use case (cf. Corollary 1 and its proof in §B.3).

$$\Sigma = \{\ \begin{aligned} \phi &\to x_1, \\ \phi &\to x_2, \\ \phi &\to x_3, \\ x_1\,x_2\,x_3\,x_5\,\upsilon &\to x_4, \\ x_1\,x_3\,x_4\,\upsilon &\to x_5, \\ x_4\,\upsilon &\to x_6, \\ x_5\,\upsilon &\to x_7\ \}. \end{aligned} \qquad \Sigma^{\hookrightarrow} = \{\ \begin{aligned} \phi &\to x_1, \\ \phi &\to x_2, \\ \phi &\to x_3, \\ \phi\,\upsilon\,x_5 &\to x_4, \\ \phi\,\upsilon\,x_4 &\to x_5, \\ \phi\,\upsilon\,x_5 &\to x_6, \\ \phi\,\upsilon\,x_4 &\to x_7\ \}. \end{aligned}$$

Figure 4: Primitive fd set Σ encoding (cf. Alg. 2) the structure of Fig. 3a and its folding $\Sigma^{\hookrightarrow}$ derived by Alg. 4.

DEF. 13. *Let Σ be set of fd's on attributes U. Then, we say that Σ is **parsimonious** if it is canonical and, for all fd's $\langle X, A\rangle \in \Sigma$ with $XA \subseteq U$, there is no $Y \subset U$ such that $Y \neq X$ and $\langle Y, A\rangle \in \Sigma$.*

COROLLARY 1. *Let Σ be an fd set defined $\Sigma \triangleq$ h-encode($\mathcal{S}$) for some complete structure $\mathcal{S}$. Then Σ can be assumed parsimonious with no loss of generality at expense of time that is polynomial in $|\Sigma| \cdot |U|$.*

PROOF. See Appendix, §B.3. □

We draw attention to the significance of Theorem 1, as it sheds light on a connection between Simon's complete structures [24] and fd sets [27]. In fact, we continue to elaborate on that connection in next section to ensure the synthesis of relational schemas with properties of our interest. We may assume given fd sets to be canonical or parsimonious if they are defined by h-encode (Corollary 1).

4. SYNTHESIS '4C'

In this section we present a technique to address Problem 2. Recall that we aim at a synthesis method to ensure the produced schema $\boldsymbol{R}$ bears some desirable properties, viz., P1′ (BCNF), P3 and P4; give up P2′. Let us then consider procedure synthesize (Alg. 3). This algorithm is essentially Bernstein's [4]. In our use case, input fd sets are defined $\Sigma \triangleq$ h-encode($\mathcal{S}$) and then may safely assumed to be parsimonious. Rather than modifying that classical algorithm, we shall achieve the properties we want for a synthesized schema by a very specific manipulation on input fd set Σ.

Algorithm 3 Schema synthesis.

```
 1: procedure SYNTHESIZE(Σ : fd set)
Require: Σ given is parsimonious, and let U := Attrs(Σ)
Ensure: Returns schema R[U] in 3NF that preserves Σ
 2:     Σ' ← apply (R4) union to Σ
 3:     R ← ∅
 4:     for all ⟨X, Z⟩ ∈ Σ' do
 5:         if there is R_k[YW] ∈ R such that X ↔ Y then
 6:             R_k ← R_k ∪ XZ
 7:         else
 8:             R_{i+1} ← XZ, with designated key X
 9:             R ← R ∪ R_{i+1}
10:     return R
```

PROPOSITION 1. *Let $\boldsymbol{R}[U]$ be a relational schema, defined $\boldsymbol{R} \triangleq$ synthesize(Σ) for some canonical fd set Σ on attributes U. Then $\boldsymbol{R}$ preserves Σ, and is in 3NF but may not be in BCNF.*

PROOF. See Appendix, §B.4. □

REMARK 2. *The connection of Proposition 1 with Theorem 1 and Corollary 1 reveals an interesting result, viz., the encoding of complete structures (i.e., derived from determinate systems of mathematical equations) followed by straightforward synthesis always leads to 3NF relational schemas. It is suggestive of the precision and (verifiable) consistency of mathematical systems, in comparison to arbitrary information systems.* □

In Proposition 1 we were not concerned with the recoverability of data. Classical versions of (Alg. 3) synthesize include an (artificial) additional step to ensure the lossless join property (cf. [1, p. 257-8]). We postpone the study of Alg. 3 w.r.t. that property to §4.2.

So far, we know that any relational schema synthesized straightforwardly from its primitive fd set $\Sigma :=$ h-encode($\mathcal{S}$) is in 3NF and is dependency-preserving. Yet, it does not give us a "good" design in the sense of Problem 2 (cf. §2.4). It fails w.r.t. P1′ (BCNF) and P3 (minimal-cardinality schema). For example, synthesize over Σ given in Fig. 4 (left) produces $|\boldsymbol{R}| = 5$, while we target at a more succinct, less decomposed schema. We shall give up strict dependency preservation to go beyond 3NF towards BCNF for a more compact representation of the causal ordering 'embedded' in Σ (cf. §4.1). For hypothesis management, a less redundant schema (BCNF over 3NF) matters not because of update anomalies but succintness. Interestingly, Arenas and Libkin have shown in information-theoretic terms how 'non-redundant' schemes in BCNF are [2].

4.1 Reflexive Pseudo-Transitive Reasoning

We seek the most succinct schema that somewhat preserves the causal ordering of the fd set given. That is achievable (see e.g., $\Sigma^{\hookrightarrow}$ in Fig. 4, right) by reflexive pseudo-transitive reasoning over Σ.

DEF. 14. *Let Σ be a set of fd's on attributes U. Then $\Sigma^{\triangleright}$, the **reflexive pseudo-transitive closure** of Σ, is the set $\Sigma^{\triangleright} \supseteq \Sigma$ such that $X \to Y$ is in $\Sigma^{\triangleright}$, with $XY \subseteq U$, iff it can be derived from a finite (possibly empty) application of rules R0, R5 over fd's in Σ. In that case, we may write $X \xrightarrow{\triangleright} Y$ and omit 'w.r.t. Σ' if it can be understood from the context.*

We are in fact interested in a very specific proper subset of $\Sigma^{\triangleright}$, say, a kernel of fd's in $\Sigma^{\triangleright}$ that gives a "compact" representation of the causal ordering 'embedded' in Σ. Note that, to characterize such special subset we shall need to be careful w.r.t. the presence of cycles in the causal ordering.

DEF. 15. *Let Σ be a set of fd's on attributes U, and $\langle X,A\rangle \in \Sigma^{\triangleright}$ with $XA \subseteq U$. We say that $X \to A$ is **folded** (w.r.t. Σ), and write $X \xrightarrow{\hookrightarrow} A$, if it is non-trivial and for no $Y \subset U$ with $Y \not\supseteq X$, we have $Y \to X$ and $X \not\to Y$ in Σ^{+}.*

The intuition of Def. 15 is that an fd is folded when there is no sense in going on with pseudo-transitive reasoning over it anymore. Given an fd $X \to A$ in fd set Σ, we shall be able to find some folded fd $Z \to A$ by applying (R5) pseudo-transitivity as much as possible while ruling out cyclic or trivial fd's in some clever way.

DEF. 16. *Let Σ be an fd set on attributes U, and $\langle X, A\rangle \in \Sigma$ be an fd with $XA \subseteq U$. Then,*

*(a) $A^{\curvearrowright}$, the **(attribute) folding** of A (w.r.t. Σ) is an attribute set $Z \subset U$ such that $Z \xrightarrow{\curvearrowright} A$;*

*(b) Accordingly, $\Sigma^{\curvearrowright}$, the **folding** of Σ, is a proper subset $\Sigma^{\curvearrowright} \subset \Sigma^{\triangleright}$ such that an fd $\langle Z, A\rangle \in \Sigma^{\triangleright}$ is in $\Sigma^{\curvearrowright}$ iff $X \xrightarrow{\curvearrowright} A$ for some $Z \subset U$.*

EXAMPLE 1. *(continued). Fig. 4 shows an fd set Σ (left) and its folding $\Sigma^{\curvearrowright}$ (right). Note that the folding can be obtained by computing the attribute folding for A in each fd $X \to A$ in Σ. We illustrate below some reasoning steps to partially compute an attribute folding.*

1.	$\phi \cup x_4$	$\to$	x_5	[consider given]
2.	$\phi \cup x_5$	$\to$	x_4	[consider given]
3.	$x_4 \cup$	$\to$	x_6	[given]
4. $\therefore$	$\phi \cup x_5$	$\to$	x_6	[R5 over (2), (3)].

Note that (4) is still amenable to further application of R5, say, over (1), (4), to derive (5) $\phi \cup x_4 \to x_6$. However, even though (4) and (5) have (resp.) the form $X \to A$ and $Y \to A$ with $Y \to X$, we have $X \to Y$ as well which characterizes a cycle. In fact, (4) itself satisfies Def. 15 and then is folded (w.r.t. Σ from Fig. 4). The same holds for (1) and (2). □

LEMMA 2. *Let Σ be a parsimonious fd set on attributes U, and $\langle X, A\rangle \in \Sigma$ be an fd with $XA \subseteq U$. Then $A^{\curvearrowright}$, the attribute folding of A (w.r.t. Σ) exists. Moreover, if Σ is parsimonious then $A^{\curvearrowright}$ is unique.*

PROOF. See Appendix, §B.5. □

We give an original algorithm (Alg. 4) to compute the folding of an fd set. At its core there lies (Alg. 5) AFolding, which can be understood as a non-obvious variant of XClosure (cf. Alg. 7) designed for acyclic reflexive pseudo-transitivity reasoning. In order to compute the folding of attribute A in fd $\langle X, A\rangle \in \Sigma$, algorithm AFolding backtraces the causal ordering 'embedded' in Σ towards A. Analogously, in terms of the directed graph G_{φ_t} induced by the causal ordering (see Fig. 3c), that would comprise graph traversal to identify the nodes x_p that have $x \mapsto A$ in their reachability, $x_p \rightsquigarrow x$. Rather, AFolding's processing of the causal ordering is fully symbolic based on Armstrong's rewrite rules R0, R5.

EXAMPLE 2. *Cyclicity in an fd set Σ may have the effect of making its folding $\Sigma^{\curvearrowright}$ to degenerate to Σ itself. For instance, consider $\Sigma = \{A \to B,\ B \to A\}$. Note that Σ is parsimonious, and AFolding (w.r.t. Σ) is B given A, and A given B. That is, $\Sigma^{\curvearrowright} = \Sigma$.* □

THEOREM 2. *Let Σ be a parsimonious fd set on attributes U, and A be an attribute with $\langle X, A\rangle \in \Sigma$ with $XA \subseteq U$. Then AFolding(Σ, A) correctly computes $A^{\curvearrowright}$, the attribute folding of A (w.r.t. Σ) in time $O(n^2)$ in $|\Sigma| \cdot |U|$.*

Algorithm 4 Folding of an fd set.

```
1: procedure FOLDING(Σ: fd set)
Require: Σ given is parsimonious
Ensure: Returns fd set Γ = Σ^↷, the folding of Σ
2:     Γ ← ∅
3:     for all ⟨X, A⟩ ∈ Σ do
4:         Z ← AFolding(Σ, A)
5:         Γ ← Γ ∪ ⟨Z, A⟩
6:     return Γ
```

Algorithm 5 Folding of an attribute w.r.t. an fd set.

```
1: procedure AFOLDING(Σ: fd set, A: attribute)
Require: Σ is parsimonious
Ensure: Returns A^↷, the attribute folding of A (w.r.t. Σ)
2:     Λ ← ∅                                   ▷ consumed attrs.
3:     Δ ← ∅                                   ▷ consumed fd's
4:     A* ← A          ▷ store "causal parent" attrs. of A
5:     size ← 0
6:     while size < |A*| do      ▷ halt when A^(i+1) = A^(i)
7:         size ← |A*|
8:         Σ ← Σ \ Δ
9:         for all ⟨Y, B⟩ ∈ Σ do
10:            if B ∈ A* then
11:                Δ ← Δ ∪ {⟨Y, B⟩}              ▷ consume fd
12:                A* ← A* ∪ Y
13:                if Y ∩ Λ = ∅ then             ▷ non-cyclic fd
14:                    Λ ← Λ ∪ B                 ▷ consume attr.
15:    return A* \ Λ
```

PROOF. For the proof roadmap, note that AFolding is monotone and terminates precisely when $A^{(i+1)} = A^{(i)}$, where $A^{(i)}$ denotes the attributes in $A^\star$ at step i of the outer loop. The folding $A^{\curvearrowright}$ of A at step i is $A^{(i)} \setminus \Lambda^{(i)}$. We shall prove by induction, given attribute A in fd $X \to A$ in parsimonious Σ, that $A^\star \setminus \Lambda$ returned by AFolding(Σ, A) is the unique attribute folding $A^{\curvearrowright}$ of A. See Appendix, §B.6. □

REMARK 3. *Beeri and Bernstein gave a straightforward optimization to (Alg. 7) XClosure to make it linear in $|\Sigma| \cdot |U|$ (cf. [3, p. 43-5]). It applies likewise to (Alg. 5) AFolding, but we omit its tedious exposure here and simply consider that AFolding can be implemented to be $O(n)$ in $|\Sigma| \cdot |U|$.*[2] □

COROLLARY 2. *Let Σ be a canonical fd set on attributes U. Then algorithm folding(Σ) correctly computes $\Sigma^{\curvearrowright}$, the folding of Σ in time that is $f(n)\,\Theta(n)$ in the size $|\Sigma| \cdot |U|$, where $f(n)$ is the time complexity of (Alg. 5) AFolding.*

PROOF. See Appendix, §B.7. □

Finally, another property of the folding of an fd set which shall be useful to know is given by Proposition 2.

PROPOSITION 2. *Let Σ be an fd set, and $\Sigma^{\curvearrowright}$ its folding. If Σ is parsimonious then so is $\Sigma^{\curvearrowright}$.*

PROOF. See Appendix, §B.8. □

[2] In short, it shall require one more auxiliary data structure to keep track, for each fd not yet consumed, of how many attributes not yet consumed appear in its rhs.

4.2 Schema Synthesis over the Folding $\Sigma^{\hookrightarrow}$

We motivate our goal of computing the folding to carry out schema synthesis over it by means of Example 3.

EXAMPLE 3. *Let us consider canonical fd set* $\Sigma = \{A \to B,\ B \to C\}$ *over attributes* $U = \{A, B, C\}$*, and a tentative schema containing a single relation* $R[ABC]$*. This relation is not in BCNF because, for one,* $B \to C$ *violates it (*$C \not\subseteq B$ *but* B *is not a superkey for* R*). A typical approach to provide a BCNF schema is to apply a 'decomposition into BCNF' algorithm (cf. [1]) to get BCNF schema* $\boldsymbol{R} = R_1[AB] \cup R_2[BC]$*. Instead, it suffices for us to consider the folding* $\Sigma^{\hookrightarrow} = \{A \to B,\ A \to C\}$ *of* Σ*. By straightforward synthesis, we generate* $R[ABC]$ *which is BCNF w.r.t.* $\Sigma^{\hookrightarrow}$*.* □

Schema synthesis over $\Sigma^{\hookrightarrow}$, the folding of a parsimonious fd set Σ, gives up preservation of Σ to target at a BCNF schema that somewhat preserves the causal ordering 'embedded' in Σ, i.e., preserves $\Sigma^{\hookrightarrow}$. Now we review the properties of relational schema $\boldsymbol{R}$ as then synthesized over the folding $\Sigma^{\hookrightarrow}$ of Σ. Recall from Proposition 1 that synthesis over Σ may render a schema $\boldsymbol{R}$ not in BCNF. The problem of deciding whether a given $\boldsymbol{R}$ is in BCNF is NP-complete [1, p. 256]. However, by Theorem 3 we shall guarantee the BCNF property a priori for every schema synthesized over the folding.

THEOREM 3. *Let* $\boldsymbol{R}[U]$ *be a relational schema, defined* $\boldsymbol{R} \triangleq$ *synthesize*$(\Sigma^{\hookrightarrow})$*, where* $\Sigma^{\hookrightarrow}$ *is the folding of parsimonious fd set* Σ *on attributes* U*. We claim that* $\boldsymbol{R}$ *is in BCNF, is minimal-cardinality and preserves* $\Sigma^{\hookrightarrow}$*.*

PROOF. See Appendix, §B.9. □

PROPOSITION 3. *Let* $\boldsymbol{R}[U] \triangleq$ *synthesize*$(\Sigma^{\hookrightarrow})$ *be a relational schema with* $|\boldsymbol{R}| \geq 2$*, where* $\Sigma^{\hookrightarrow}$ *is the folding of a parsimonious fd set* $\Sigma \triangleq$ *h-encode*$(\mathcal{S})$ *on attributes* U*. Then* $\boldsymbol{R}$ *has a lossless join (w.r.t.* $\Sigma^{\hookrightarrow}$*) iff, for all* $R_i[XZ] \in \boldsymbol{R}$ *with key constraint* $X \to Z$*, we have* $X \to U$ *or there is* $R_j[YW] \in \boldsymbol{R}$ *such that* $X \subset Y$*.*

PROOF. See Appendix, §B.10. □

REMARK 4. *An alternative approach (cf. [27, p. 411]) to ensure the lossless join property w.r.t.* $\Sigma^{\hookrightarrow}$ *over attributes* U *is to render an additional "artificial" scheme* $R_{i+1}[X]$*, where* X *is any superkey for* U*, in order to get* $\boldsymbol{R}' := \boldsymbol{R} \cup R_{i+1}[X]$*. Such* $\boldsymbol{R}'$ *is in BCNF and has a lossless join for sure but is not the minimal-cardinality schema in BCNF and then is not considered here.* □

EXAMPLE 4. *Apply* $\boldsymbol{R} \triangleq$ *synthesize*$(\Sigma^{\hookrightarrow})$*, where* $\Sigma^{\hookrightarrow}$ *is given in Fig. 4 (right). Then we get* $\boldsymbol{R} = \{R_1[\phi\, x_1\, x_2\, x_3],\ R_2[\phi\, \upsilon\, x_5\, x_4\, x_6\, x_7]\}$*, which is in BCNF, preserves* $\Sigma^{\hookrightarrow}$ *and has a lossless join. Now, let us take a slightly different fd set* $\Gamma \triangleq \Sigma \cup \{x_1\, x_9\, \upsilon \to x_8,\ \ x_2\, x_8\, \upsilon \to x_9\}$*. By applying* $\boldsymbol{R}' \triangleq$ *synthesize*$(\Gamma^{\hookrightarrow})$*, we get* $\boldsymbol{R}' = \boldsymbol{R} \cup \{R_3[\phi\, \upsilon\, x_8, x_9]\}$*, which is in BCNF, preserves* $\Sigma^{\hookrightarrow}$ *but does not have a lossless join. It turns out that* $\Gamma^{\hookrightarrow}$ *"embeds" two subsets of strongly coupled variables (attributes), viz.* $\{x_4, x_5, x_6, x_7\}$ *and* $\{x_8, x_9\}$ *that are not "causally connected" to each other.* □

CONJECTURE 1. *The lossless join property is reducible to the structure* $\mathcal{S}$ *given as input to the pipeline.*

We comment on Conjecture 1 in some detail in §C.1. In the converse direction, we bring in Def. 17 SEM's concepts into data dependency language.

DEF. 17. *Let* $H[XZ]$ *be a relation with key constraint* $X \to Z$*. We say that* $X \to Z$ *is a* ***φ-fd*** *over* ***exogenous*** *attributes* Z *(and exogenous relation* $H[XZ]$*) if* $\upsilon \notin X$*. We say that it is an* ***υ-fd*** *over* ***endogenous*** *attributes* Z *(and endogenous relation* $H[XZ]$*) otherwise.*

5. SYNTHESIS '4U'

In this section we present a technique to address Problem 3. At this stage of the pipeline, relational schema $\boldsymbol{H}$ has been synthesized and datasets computed from the hypotheses under alternative trials (input settings) are loaded into it. The challenge now is how to render the U-relations $\boldsymbol{Y}$. Before proceeding, we consider Example 5, which is admittedly small but fairly representative to illustrate how to deal with correlations in the predictive data of deterministic hypotheses.

EXAMPLE 5. *We explore three slightly different theoretical models in population dynamics with applications in Ecology, Epidemics, Economics, etc: (2) Malthus' model, (3)the logistic equation and (4) the Lotka-Volterra model. In practice, such equations are meant to be extracted from* MathML*-compliant XML files (cf. §6.2). For now, consider that the ordinary differential equation notation '*$\dot{x}$*' is read 'variable* x *is a function of time* t *given initial condition* x_0*.'*

$$\dot{x} = rx \tag{2}$$

$$\dot{x} = r(C - x)x \tag{3}$$

$$\begin{cases} \dot{x} &= x(b - py) \\ \dot{y} &= y(rx - d) \end{cases} \tag{4}$$

The models are completed (by the user) with additional equations to provide the values of exogenous variables (or "input parameters"),[3] *e.g.,* $x_0 = 200$*,* $r = 10$*, such that we have structures* $\mathcal{S}_k(\mathcal{E}_k, \mathcal{V}_k)$*, for* $k = 1..3$*,*

- $\mathcal{E}_1 = \{\, f_1(t),\ f_2(x_0),\ f_3(r),\ f_4(x, t, x_0, r)\,\}$*;*
- $\mathcal{E}_2 = \{f_1(t),\ f_2(x_0),\ f_3(C),\ f_4(r),\ f_5(x, t, x_0, C, r)\}$*;*
- $\mathcal{E}_3 = \{\, f_1(t),\ f_2(x_0),\ f_3(b),\ f_4(p),\ f_5(y_0),\ f_6(d),\ f_7(r),\ f_8(x, t, x_0, b, p, y),\ f_9(y, t, y_0, d, r, x)\,\}$*.*

Fig. 5 shows the fd sets encoded from structures $\mathcal{S}_k$ *above.*[4] *We shall also consider trial datasets for hypothesis* $\upsilon = 3$ *(viz., the Lotka-Volterra model), which are loaded into the synthesized (certain) schemes in* $\boldsymbol{H}_3$ *as shown in Fig. 5. Note that the fd's in* Σ_3 *are violated by relations* H_3^1*,* H_3^2*, but we admit a special attribute 'trial id'* tid *into their key constraints for a trivial repair (provisionally, yet at the* ETL *stage of the pipeline) until uncertainty is introduced in a controlled way by synthesis '4U' (*U-intro *stage, cf. Fig. 2).* □

Given *certain* relations $\boldsymbol{H}$, synthesis '4U' has two parts: process the uncertainty of exogenous relations (*u-factorization*) and of endogenous relations (*u-propagation*).

[3]Given $\mathcal{S}(\mathcal{E}, \mathcal{V})$, it is actually a task of the encoding algorithm (viz., COA_t's) to infer whether its variables $x \in \mathcal{V}$ are exogenous or endogenous by processing its causal ordering.

[4]Domain variables like time t require a special treatment by h-encode to suppress an fd $\phi \to t$. This is coped with by providing it an additional argument ℓ informing that the $(\ell \times \ell)$-first block of matrix A_S is kept for domain variables.

$\Sigma_1 = \{\ \phi \to x_0,\ \phi \to r,\ x_0\, r\, t\, \upsilon \to x\ \}.$

$\Sigma_2 = \{\ \phi \to x_0,\ \phi \to C,\ \phi \to r,\ x_0\, C\, r\, t\, \upsilon \to x\ \}.$

$\Sigma_3 = \{\ \phi \to x_0,\ \phi \to b,\ \phi \to p,\ \phi \to y_0,\ \phi \to d,\ \phi \to r,\ x_0\, b\, p\, t\, \upsilon\, y \to x,\ y_0\, d\, r\, t\, \upsilon\, x \to y\ \}.$

H_3^1

tid	ϕ	x_0	b	p	y_0	d	r
1	1	3	1	1	6	1	1
2	1	10	1.5	1	5	3	1
3	1	30	.5	.02	4	.5	.02
4	1	30	.4	.02	4	.8	.02
5	1	30	.4	.018	4	.8	.023
6	1	30	.397	.018	4	.786	.023

$\bowtie$

H_3^2

tid	ϕ	υ	t	x	y
1	1	3	0	3	6
1	1	3	..	..	..
..	1	3	..	..	..
6	1	3	0	30	4
6	1	3	5	50.1	62.9
6	1	3	10	13.8	8.65
6	1	3	15	79.3	8.23
6	1	3	20	12.6	30.7
6	1	3	..	..	..

Figure 5: Resources from Example 5. (Left). Primitive fd sets extracted from the given structures $\mathcal{S}_k(\mathcal{E}_k, \mathcal{V}_k)$ for $k = 1..3$. (Right). *Certain* relations $\boldsymbol{H}_3 = \{H_3^1,\ H_3^2\}$ of hypothesis $\upsilon = 3$ loaded with trial datasets identified by special attribute tid.

5.1 U-Factorization

As we have seen in §2.1, the repair-key operation allows one to create a discrete random variable in order to repair an argument key in a given relation. Our goal here is to devise a technique to perform such operation in a principled way for hypothesis management. It is a basic design principle to have exactly one random variable for each distinct uncertainty factor ('u-factor' for short), which requires carefully identifying the actual sources of uncertainty present in relations $\boldsymbol{H}$.

The multiplicity of (competing) hypotheses is itself a standard one, viz., the *theoretical* u-factor. Consider an 'explanation' table like H_0 in Fig. 1, which stores (as foreign keys) all hypotheses available and their target phenomena. We can take such H_0 as explanation table for the three hypotheses of Example 5. Then a discrete random variable x_0 (not to be confused with variable a $x_0 \in \mathcal{V}$) is defined into $Y_0[\, V_0 D_0 \,|\, \phi\, \upsilon\,]$ by query formula (1). U-relation Y_0 is considered standard in synthesis '4U,' as the repair of ϕ as a key in (standard) H_0.

Hypotheses, though, are (abstract) 'universal statements' [18]. In order to produce a (concrete) valuation over their endogenous attributes (predictions), one has to inquire into some particular 'situated' phenomenon ϕ and tentatively assign a valuation over the exogenous attributes, which can be eventually tuned for a target ϕ. The multiplicity of such (competing) empirical estimations for a hypothesis k leads to Problem 4, viz., learning *empirical* u-factors for each $\boldsymbol{H}_k \subseteq \boldsymbol{H}$.

Problem 4. *Let $H_k^\ell[XZ] \in \boldsymbol{H}_k$ be an exogenous relation with key constraint $X \to Z$. Once H_k^ℓ is loaded with trial data, the problem of* ***u-factor learning*** *is:*

1. *to* ***infer*** *in H_k^ℓ "casual" fd's $B_i \leftrightarrow B_j \notin \Sigma_k$ (strong input correlations), where B_i, $B_j \in Z$;*
2. *to form maximal* ***groups*** *G_1, ..., $G_n \subseteq Z$ of attributes such that for all B_i, $B_j \in G_a$, the casual fd's $B_i \leftrightarrow B_j$ hold in H_k^ℓ;*
3. *to pick, for each group G_a, any $A \in G_a$ as a* ***pivot*** *representative and insert $A \to B$ into an fd set Γ_k for all $B \in (G_a \setminus A)$.*

Problem 4 is dominated by the (problem of) discovery of fd's in a relation, which is not really a new problem (e.g., see [14]). We then keep focus on the synthesis '4U' as a whole and omit our detailed u-factor-learning algorithm in particular. Its output, fd set Γ_k, is then filled in (completed) with the υ-fd's from Σ_k.

$\Gamma_3 = \{\ x_0\, b\, p\, t\, \upsilon\, y \to x,\ y_0\, d\, r\, t\, \upsilon\, x \to y\ \}.$

$\Gamma_3^{\looparrowright} = \{\ x_0\, b\, p\, t\, \upsilon\, y \to x,\ x_0\, b\, p\, t\, \upsilon\, x \to y\ \}.$

Figure 6: Fd set Γ_3 (compare with Σ_3) and its folding $\Gamma_3^{\looparrowright}$.

For illustration consider hypothesis $\upsilon = 3$ and its trial input data recorded in H_3^1 in Fig. 5. We show its corresponding fd set Γ_3 in Fig. 6 (left). Recall that, as a result of synthesis '4C,' relation H_3^1 is in BCNF w.r.t. $\Sigma_3^{\looparrowright}$. Since its attributes have been inferred exogenous in the given hypothesis (cf. Proposition 5), they are then officially unrelated. In fact, by "casual" fd's we mean correlations that, for a set of experimental trials, may occasionally show up in the trial input data — e.g., $x_0 \leftrightarrow y_0$ hold in H_3^1, but not because x_0 and y_0 are related in principle (theory).

Once fd set Γ_k is output by u-factor learning, its folding $\Gamma_k^{\looparrowright}$ shall be given with $\boldsymbol{H}_k$ as input to accomplish u-factorization for hypothesis k algorithmically. We shall employ a notion of u-factor decomposition formulated in Def. 18 into query formula (5).

Def. 18. *Let $R[A_pW] \in \boldsymbol{R}$ be an exogenous scheme with $\boldsymbol{R} \triangleq$ synthesize(Φ) designed over subset Φ of ϕ-fd's in $\Gamma_k^{\looparrowright}$; and let $H_k^\ell[X_\ell L] \in \boldsymbol{H}_k$ be an exogenous relation with (violated) key constraint $\langle X_\ell, L\rangle \in \Sigma_k^{\looparrowright}$ with $A_pW \subseteq L$. Then the exogenous U-relation $Y_k^p[V_pD_p \,|\, X_\ell A_p]$ for sketched scheme $R[A_pW]$ is defined by query formula (5) in p-WSA's extension of relational algebra,*

$$Y_k^p := \pi_{X_\ell A_p}(\textsf{repair-key}_{X_\ell @ \mathsf{count}}(\gamma_{X_\ell A_p,\, \mathsf{count}(*)}(H_k^\ell))) \quad (5)$$

where γ is relational algebra's grouping operator. Let $G_a = A_pW$. We say that Y_k^p is a ***u-factor projection*** *of $H_k^\ell[X_\ell L]$ if $A_p \leftrightarrow B$ hold in H_k^ℓ for all $B \in G_a$ and for no $C \in (L \setminus G_a)$ we have $C \to B$ or $B \to C$.*

Proposition 4. *Let $H_k^\ell[X_\ell L] \in \boldsymbol{H}_k$ be an exogenous relation with (violated) key constraint $\langle X_\ell, L\rangle \in \Sigma_k^{\looparrowright}$. Then,*

(a) for any pair $Y_k^i[\, V_iD_i \,|\, X_\ell A_i]$, $Y_k^j[\, V_jD_j \,|\, X_\ell A_j]$ of u-factor projections of H_k^ℓ, they are independent.

(b) the join $\bowtie_{i=1}^m Y_k^i[\, V_iD_i \,|\, X_\ell A_i\,]$ of all u-factor projections of H_k^ℓ is lossless w.r.t. $\pi_{X_\ell A_1\, A_2\, \ldots\, A_m}(\Sigma_k^{\looparrowright})$.

PROOF. See Appendix, §B.11. □

Proposition 4 is significant as it ensures all the empirical uncertainty implicit in an exogenous relation can be decomposed into u-factor projections that are (a) in fact independent, to do justice to the term 'factors,' and (b) and can be fully recovered by a lossless join. The u-factorization procedure given fd set $\Gamma_k^{\leftarrow}$ and relations $\boldsymbol{H}_k$ is described by (Alg. 6) synthesize4u (Part I).

5.2 U-Propagation

For hypothesis $\upsilon = k$, take some endogenous attribute B, and then by Def. 17 and the parsimony assumption we must have exactly one υ-fd $S \to B$ in $\Gamma_k^{\leftarrow}$. Then note that the u-factors with 'incidence' on B are in S. In comparison with synthesis '4C,' we have just unfolded B's "causal chain" out of its compact form in $\Sigma_k^{\leftarrow}$ and re-folded it fine-grained (over u-factor pivots) into $\Gamma_k^{\leftarrow}$.

Note that the υ-fd's in $\Gamma_k^{\leftarrow}$, of form $S \to B$, are meant for u-propagation. Each pivot attribute $A_j \in S$ shall be used as a surrogate to its associated random variable x_j from exogenous U-relation $Y_k^j[V_j D_j \,|\, X_\ell A_j]$ to propagate uncertainty properly into endogenous U-relations $Y_k^r[\overline{V_i D_i} \,|\, ZT]$, for $B \in T$ and each $A_j \in S \setminus Z$. This intuition is abstracted into a general p-WSA query formula (6) as given in Def. 19, and employed in (Alg. 6) synthesize4u to accomplish u-propagation (Part II).

DEF. 19. *Let $R[ST] \in \boldsymbol{R}$ be an endogenous scheme, designed $\boldsymbol{R} \triangleq$ synthesize(Υ) over subset Υ of υ-fd's in $\Gamma_k^{\leftarrow}$; and let $H_k^q[Z_q V] \in \boldsymbol{H}_k$ be an endogenous relation with (violated) key constraint $\langle Z_q, V\rangle \in \Sigma_k^{\leftarrow}$ such that $V \supseteq T$. Then the endogenous U-relation $Y_k^q[\overline{V_i D_i} \,|\, Z_q T]$ for sketched scheme $R[ST]$ is defined by formula (6),*

$$Y_k^r := \pi_{Z_q T}(\, \sigma_{\upsilon=k}(Y_0) \bowtie \pi_{\mathsf{tid},X}(J) \bowtie \pi_{\mathsf{tid},Z_q T}(H_k^q)\,) \quad (6)$$

where J is a join sub-query defined over mapping M from exogenous relations to sets of exogenous U-relations:

(a) for $H_k^\ell \in \boldsymbol{H}_k$, we have $H_k^\ell[X_\ell L] \in \mathsf{M}$ iff $L \cap S \neq \varnothing$;

(b) we take $X = \bigcup_{H_k^\ell \in \mathsf{M}} X_\ell$;

(c) we have $Y_k^j[V_j D_j | X_\ell A_j] \in \mathsf{M}(H_k^\ell)$ iff Y_k^j is a u-factor projection of H_k^ℓ with $A_j \in (L \cap S)$.

*We say that Y_k^r is a **predictive projection** of H_k^q.*

THEOREM 4. *Let $H_k^q[Z_q V] \in \boldsymbol{H}_k$ be an endogenous relation with (violated) key constraint $\langle Z_q, V\rangle \in \Sigma_k^{\leftarrow}$, and $Y_k^r[\overline{V_i D_i} \,|\, Z_q T]$ be a predictive projection of H_k^q w.r.t. $\Gamma_k^{\leftarrow}$ defined by formula (6) with $V \supseteq T$. We claim that Y_k^r correctly captures all the uncertainty present in H_k^q w.r.t. $\Gamma_k^{\leftarrow}$.*

PROOF. See Appendix, §B.12. □

Fig. 7 shows the rendered U-relations for hypothesis $\upsilon = 3$ whose relations are shown in Fig. 5. Note that $\mathsf{tid} = 6$ in H_3^2 (Fig. 5) corresponds now to $\theta = \{\mathsf{x}_0 \mapsto 3, \mathsf{x}_1 \mapsto 3, \mathsf{x}_2 \mapsto 5, \mathsf{x}_3 \mapsto 3\}$, where θ defines a particular world in $\boldsymbol{W}$ whose probability is $\mathsf{Pr}(\theta) \approx .012$. This value is derived from the marginal probabilities stored in world table W (e.g., see Fig. 1) as a result of the application of formulas (1) and (5).

Algorithm 6 Synthesis '4U' applied over folding fd set.

procedure SYNTHESIZE4U($\Gamma_k^{\leftarrow}$: fd set, $\boldsymbol{H}_k$: DB)
Require: $\Gamma_k^{\leftarrow}$ is the folding of parsimonious fd set Γ_k
Ensure: U-relational DB $\boldsymbol{Y}_k$ returned is $\boldsymbol{H}_k$ after U-intro
  $\Phi \leftarrow \varnothing$, $\Upsilon \leftarrow \varnothing$
  for all $\langle X, B\rangle \in \Gamma_k^{\leftarrow}$ **do**
    if $\upsilon \notin X$ **then**
      $\Phi \leftarrow \Phi \cup \langle X, B\rangle$ ▷ ϕ-fd over exogenous B
    else
      $\Upsilon \leftarrow \Upsilon \cup \langle X, B\rangle$ ▷ υ-fd over endogenous B

Part I: **U-factorization**
  $\mathsf{M} \leftarrow \varnothing$ ▷ store u-factor projection mappings
  $\boldsymbol{R} \leftarrow$ synthesize(Φ) ▷ design BCNF exog. schemes
  for all $R[A_p W] \in \boldsymbol{R}$ **do**
    find exog. $H_k^\ell[X_\ell L] \in \boldsymbol{H}_k$ such that $A_p \in L$
    $Y_k^i \leftarrow \pi_{X_\ell A_p}(\mathsf{repair\text{-}key}_{X_\ell @ count}(\gamma_{X_\ell A_p, count(*)}(H_k^\ell)))$
    $\boldsymbol{Y}_k \leftarrow \boldsymbol{Y}_k \cup Y_k^i$
    if $H_k^\ell \in \mathsf{M}$ **then** ▷ save mapping for further ref.
      $\mathsf{M}(H_k^\ell) \leftarrow \mathsf{M}(H_k^\ell) \cup \{Y_k^i\}$
    else
      $\mathsf{M} \leftarrow \mathsf{M} \cup \langle H_k^\ell, \{Y_k^i\}\rangle$

Part II: **U-propagation**
  $\boldsymbol{R} \leftarrow$ synthesize(Υ) ▷ design BCNF endog. schemes
  for all $R[ST] \in \boldsymbol{R}$ **do**
    $X \leftarrow \varnothing$, $J \leftarrow \varnothing$ ▷ prepare for join sub-query
    for all $H_k^\ell[X_\ell\, L] \in \mathsf{M}$ **do**
      if $L \cap S \neq \varnothing$ **then**
        $J \leftarrow J \bowtie H_k^\ell$
        $X \leftarrow X \cup X_\ell$
        for all $Y_k^i[X_\ell A_p] \in \mathsf{M}(H_k^\ell)$ **do**
          if $A_p \in S$ **then** ▷ Y_k^i is a u-factor for R
            $J \leftarrow J \bowtie Y_k^i$
    find $H_k^q[Z_q V] \in \boldsymbol{H}_k$ such that $V \supseteq T$
    $Y_k^r \leftarrow \pi_{Z_q T}(\, \sigma_{\upsilon=k}(Y_0) \bowtie \pi_{\mathsf{tid},X}(J) \bowtie \pi_{\mathsf{tid},Z_q T}(H_k^q)\,)$
    $\boldsymbol{Y}_k \leftarrow \boldsymbol{Y}_k \cup Y_k^r$
  return $\boldsymbol{Y}_k$

REMARK 5. *Observe that, although (Alg. 6) synthesize4u operates locally for each hypothesis k, the effects of synthesis '4U' (U-intro) in the pipeline are global on account of the (global) 'explanation' relation H_0 (then U-relation Y_0); e.g., see Fig. 7. In fact, the probability of each tuple, say, in endogenous U-relation Y_k^q with $\phi = p$ for hypothesis $\upsilon = k$, is distributed among all the hypotheses $\ell \neq k$ that are keyed in Y_0 under $\phi = p$.* □

In sum, the synthesis '4U' technique completes the pipeline (Fig. 2), except for the problem of conditioning (cf. [10]) which is not covered in this paper.

6. DISCUSSION

In this section we discuss related work and the applicability of our framework.

6.1 Related Work

Models and data. Haas et al. [12] propose a long-term *models-and-data* research program to address data management for *deep* predictive analytics. They discusss strategies to extend query engines for model execution within a (p-)DB. Along these lines, query optimization is understood as a more general problem

Y_3^1	$V \mapsto D$	ϕ	A_1
	$x_1 \mapsto 1$	1	3
	$x_1 \mapsto 2$	1	10
	$x_1 \mapsto 3$	1	30

Y_3^2	$V \mapsto D$	ϕ	A_2
	$x_2 \mapsto 1$	1	1
	$x_2 \mapsto 2$	1	1.5
	$x_2 \mapsto 3$	1	.5
	$x_2 \mapsto 4$	1	.4
	$x_2 \mapsto 5$	1	.397

Y_3^3	$V \mapsto D$	ϕ	A_3
	$x_3 \mapsto 1$	1	1
	$x_3 \mapsto 2$	1	.02
	$x_3 \mapsto 3$	1	.018

Y_3^4	$V_0 \mapsto D_0$	$V_1 \mapsto D_1$	$V_2 \mapsto D_2$	$V_3 \mapsto D_3$	ϕ	v	t	x	y
	$x_0 \mapsto 3$	$x_1 \mapsto 1$	$x_2 \mapsto 1$	$x_3 \mapsto 1$	1	3	0	3	6
	$x_0 \mapsto 3$	$x_1 \mapsto 1$	$x_2 \mapsto 1$	$x_3 \mapsto 1$	1	3	...	...	...
	...	...	...	...	1	3	...	...	...
	$x_0 \mapsto 3$	$x_1 \mapsto 3$	$x_2 \mapsto 5$	$x_3 \mapsto 3$	1	3	0	30	4
	$x_0 \mapsto 3$	$x_1 \mapsto 3$	$x_2 \mapsto 5$	$x_3 \mapsto 3$	1	3	5	50.1	62.9
	$x_0 \mapsto 3$	$x_1 \mapsto 3$	$x_2 \mapsto 5$	$x_3 \mapsto 3$	1	3	10	13.8	8.65
	$x_0 \mapsto 3$	$x_1 \mapsto 3$	$x_2 \mapsto 5$	$x_3 \mapsto 3$	1	3	15	79.3	8.23
	$x_0 \mapsto 3$	$x_1 \mapsto 3$	$x_2 \mapsto 5$	$x_3 \mapsto 3$	1	3	20	12.6	30.7
	$x_0 \mapsto 3$	$x_1 \mapsto 3$	$x_2 \mapsto 5$	$x_3 \mapsto 3$	1	3	..	..	..

Figure 7: U-relations rendered by synthesis '4U' for hypothesis $v = 3$ from Example 5.

with connections to algebraic solvers. Our framework in turn essentially comprises an abstraction of *hypotheses as data* [10]. It can be understood in comparison as putting models strictly into a (flattened) data perspective. For this reason it is directly applicable by building upon recent work on p-DBs [26].

Design by synthesis. Classical design by synthesis [4] was once criticized due to its too strong 'uniqueness' of fd's assumption [9, p. 443], as it reduces the problem of design to symbolic reasoning on fd's arguably neglecting semantic issues. For hypothesis management, however, design by synthesis is clearly feasible (see also §6.2), as it translates seamlessly to data dependencies the reduction made by the user herself into a tentative formal model for the studied phenomenon. In fact, synthesis methods may be as fruitful for p-DB design as they are for GM design (cf. [8, 6]).

Reasoning over fd's. The concept of fd folding and design of (Alg. 5) AFolding as a variant of XClosure, is (arguably) a non-obvious approach to the problem of processing the causal ordering of a hypothesis via acyclic reflexive pseudo-transitive reasoning. To the best of our knowledge, such a specific form of fd reasoning over fd's was an yet unexplored problem in the database research literature.

Causality in AI. The work of Dash and Druzdzel (cf. [7]) is peculiar in that it explores SEMs to model deterministic (quantitative) hypotheses into GMs (like in Simon's early work), while a majority of work in their field is devoted to model statistical (qualitative) hypotheses [22, 6]. In comparison, we are applying SEMs to model deterministic hypotheses into U-relational DBs.

Causality in DBs. Our encoding of equations into fd's (constraints/correlations) captures the causal chain from exogenous (input) to endogenous (output) tuples. Fd's are rich, stronger information that can be exploited in reasoning about causality in a DB for the sake of explanation and sensitivity analysis. In comparison with Meliou et. al [20] and Kanagal et. al [16], we are processing causality at schema level. To our knowledge this is the first work to address causal reasoning in the presence of constraints (viz., fd's), called for as worth of future work by Meliou et al. [20, p. 3].

Hypothesis encoding. Finally, our framework is comparable with Bioinformatics' initiatives that address hypothesis encoding into the RDF data model [25]. We point out the Robot Scientist, HyBrow and SWAN (cf. [25]), all of which consist in some ad-hoc RDF encoding of sequence and genome analysis hypotheses under varying levels of structure (viz., from 'gene G has function A' statements to free text). Our framework in turn consists in the U-relational encoding of hypotheses from mathematical equations, which is (to our knowledge) the first work on hypothesis relational encoding.

6.2 Applicability

Essentially, the formal framework proposed in this paper is meant for (large-scale) hypothesis management and predictive analytics directly in the data under support of the querying capabilities of a p-DB.

Realistic assumptions. The core assumption of our framework is that the hypotheses are given in a formal specification which is encodable into a SEM that is complete (satisfies Def. 8). Also, as a semantic assumption which is standard in scientific modeling, we consider a one-to-one correspondence between real-world entities and variable/attribute symbols within a structure, and that all of them must appear in some of its equations/fd's. For most science use cases involving deterministic models (if not all), such assumptions are quite reasonable. It can be a topic of future work to explore business use cases as well.

Hypothesis learning. The (user) method for hypothesis formation is irrelevant to our framework, as long as the resulting hypothesis is encodable into a SEM. So, a promising use case is to incorporate machine learning methods into our framework to scale up the formation/extraction of hypotheses and evaluate them under the querying capabilities of a p-DB. Consider, e.g., learning the equations, say, from Eureqa [23].[5]

Qualitative hypotheses. Although the method is primarely motivated by computational science (usually involving differential equations), in fact it is applicable to qualitative deterministic models as well. Boolean Networks, e.g., consist in sets of functions $f(x_1, x_2, .., x_n)$, where f is a Boolean expression. Several kinds of dynamical system can be modeled in this formalism. Applications have grown out of gene regulatory network to social network and stock market predictive analytics. Even if richer semantics is considered (e.g., fuzzy logics), our encoding method is applicable likewise, as long as the equations are still deterministic.

Model repositories. Recent initiatives have been fostering large-scale model integration, sharing and reproducibility in the computational sciences (e.g., [13, 5, 15]). They are growing reasonably fast on the web, (i) promoting some MathML-based schema as a standard for model specification, but (ii) with limited integrity and lack of support for rating/ranking competing models. For those two reasons, they provide a strong use case for our method of hypothesis management. The Physiome project [15] , e.g., is planned to integrate very large deterministic models of human physiology. A fairly simple model of the human cardiovascular system has about 630+ variables (or equations, as $|\mathcal{E}| = |\mathcal{V}|$).

Initial experiments. We have run an applicability study comprising the whole pipeline of Fig. 2 in a realistic use case scenario extracted from the Physiome project [11]. Our initial experiments have shown that the pipeline can in fact be processed efficiently. All

[5] `http://creativemachines.cornell.edu/Eureqa`.

hypotheses extracted and analyzed happened to satisfy the lossless join property (cf. Appendix, C.1).

7. CONCLUSIONS

In this paper we have presented specific technical developments over the Υ-DB vision [10]. In short, we have shown how to encode deterministic hypotheses as uncertain data, viz., as constraints and correlations into U-relational DBs.

Although the pipeline (Fig. 2) is motivated by a very concrete class of applications, it raised some non-trivial theoretical issues and required a new design-theoretic framework in view of a principled DB research solution — not only to the specific technical problems P1-P3 (cf. §2.4) but to the pipeline as a whole in view of enabling hypothesis management and for predictive analytics.

This work can be understood as revisiting and making effective use of classical design theory in a modern context to address new problems.

8. ACKNOWLEDGMENTS

This research has been supported by the Brazilian funding agencies CNPq (grants n^o 141838/2011-6, 309494/2012-5) and FAPERJ (grants INCT-MACC E-26/170.030/2008, 'Nota 10' E-26/100.286/2013). We thank IBM for a Ph.D. Fellowship 2013-2014.

9. REFERENCES

[1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.

[2] M. Arenas and L. Libkin. An information-theoretic approach to normal forms for relational and xml data. *J. ACM*, 52(2):246–283, 2005.

[3] C. Beeri and P. Bernstein. Computational problems related to the design of normal form relational schemas. *ACM TODS*, 4(1):30–59, 1979.

[4] P. Bernstein. Synthesizing third normal form relations from functional dependencies. *ACM Trans. on Database Systems*, 1(4):277–98, 1976.

[5] V. Chelliah, C. Laibe, and N. Le Novère. BioModels Database: A repository of mathematical models of biological processes. *Method. Mol. Biol.*, (1021):189–99, 2013.

[6] A. Darwiche. Bayesian networks. *Comm. ACM*, 53(12):80–90, 2010.

[7] D. Dash and M. J. Druzdzel. A note on the correctness of the causal ordering algorithm. *Artif. Intell.*, 172(15):1800–8, 2008.

[8] M. J. Druzdzel and H. Simon. Causality in Bayesian belief networks. In *Proc. of the 9th International Conf. on Uncertainty in Artificial Intelligence*, pages 3–11, 1993.

[9] R. Fagin. The decomposition versus synthetic approach to relational database design. In *VLDB*, volume 3, pages 441–6, 1977.

[10] B. Gonçalves and F. Porto. Υ-DB: Managing scientific hypotheses as uncertain data. *PVLDB*, 7(11):959–62, 2014.

[11] B. Gonçalves and F. Porto. Managing large-scale scientific hypotheses as uncertain data with support for predictive analytics. Technical report, LNCC, 2014. `http://arxiv.org/abs/1405.5905`.

[12] P. Haas, P. Maglio, P. Selinger, and W. Tan. Data is dead... without what-if models. *PVLDB*, 4(12):1486–9, 2011.

[13] M. Hines, T. Morse, M. Migliore, N. Carnevale, and G. Shepherd. ModelDB: A database to support computational neuroscience. *J Comput Neurosci*, 17(1):7–11, 2004.

[14] Y. Huhtala et al. TANE: An efficient algorithm for discovering functional and approximate dependencies. *Comput. J.*, 42(2):100–11, 1999.

[15] P. J. Hunter and T. K. Borg. Integration from proteins to organs: the Physiome Project. *Nat Rev Mol Cell Biol.*, 4(3):237–43, 2003.

[16] B. Kanagal, J. Li, and A. Deshpande. Sensitivity analysis and explanations for robust query evaluation in probabilistic databases. In *SIGMOD*, pages 841–52, 2011.

[17] C. Koch. *MayBMS: A system for managing large uncertain and probabilistic databases*. In C. Aggarwal (ed.), Managing and Mining Uncertain Data, Chapter 6. Springer-Verlag, 2009.

[18] J. Losee. *A historical introduction to the philosophy of science*. Oxford Univ. Press, 4th ed., 2001.

[19] D. Maier. *Theory of Relational Databases*. Computer Science Press, 1983.

[20] A. Meliou et al. Causality in databases. *IEEE Data Eng. Bull.*, 33(3):59–67, 2010.

[21] S. L. Osborn. Testing for existance of a covering Boyce-Codd normal form. *Information Processing Letters*, 8(1):11–4, 1979.

[22] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge Univ. Press, 2000.

[23] M. Schmidt and H. Lipson. Distilling free-form natural laws from experimental data. *Science*, 324(5923):81–5, 2009.

[24] H. Simon. *Causal ordering and identifiability*. In Hood & Koopmans (eds.), Studies in Econometric Methods, Chapter 3, John Wiley & Sons, 1953.

[25] L. Soldatova and A. Rzhetsky. Representation of research hypotheses. *J Biomed Sem*, 2(S2), 2011.

[26] D. Suciu, D. Olteanu, C. Ré, and C. Koch. *Probabilistic Databases*. Morgan & Claypool Publishers, 2011.

[27] J. Ullman. *Principles of Databases and Knowledge-Base Systems*. Computer Science Press, 1988.

APPENDIX

A. AUXILIARY ALGORITHMS

We list some reference algorithms from the literature.

Algorithm 7 Attribute closure X^+ (cf. [27, p. 388]).

```
 1: procedure XClosure(Σ: fd set, X: attribute set)
Require: Σ is an fd set, X is a non-empty attribute set
Ensure: X+ is the attribute closure of X w.r.t. Σ
 2:     size ← 0
 3:     X+ ← X
 4:     Γ ← Σ
 5:     while size < |X+| do
 6:         size ← |X+|
 7:         for all ⟨Y, Z⟩ ∈ Γ do
 8:             if Y ⊆ X+ then
 9:                 X+ ← X+ ∪ Z
10:                 Γ ← Γ \ ⟨Y, Z⟩
11:     return X+
```

Algorithm 8 Left-reduced cover for a given fd set [19].

```
 1: procedure LEFT-REDUCE(Σ: fd set)
Require: Σ is an fd set
Ensure: Γ is a left-reduced cover for Σ
 2:     Γ ← Σ
 3:     for all ⟨X, Y⟩ ∈ Γ do
 4:         for all A ∈ X do
 5:             if member(Σ, ⟨X\A, Y⟩) then
 6:                 Σ ← Σ \ ⟨X, Y⟩ ∪ ⟨X\A, Y⟩
 7:     return Σ
```

Algorithm 9 Membership in the closure of an fd set [27].

```
 1: procedure MEMBER(Σ: fd set, ⟨X, Y⟩: fd)
Require: Σ is an fd set, ⟨X, Y⟩ is an fd
Ensure: A decision is returned
 2:     if Y ⊆ XClosure(Σ, X) then
 3:         return yes
 4:     else
 5:         return no
```

B. DETAILED PROOFS

B.1 Proof of Lemma 1

"Let Σ be a (Def. 6-a) singleton-rhs fd set on attributes U. Then $\langle X, A\rangle \in \Sigma^+$ with $XA \subseteq U$ only if $A \subseteq X$ or there is non-trivial $\langle Y, A\rangle \in \Sigma$ for some $Y \subset U$."

Proof. By Lemma 3 (below), we know that $X \to A \in \Sigma^+$ iff $A \subseteq X^+$. We need to prove that if $A \nsubseteq X$ and there is no $Y \to A$ in singleton-rhs Σ, then $A \nsubseteq X^+$. But this is equivalent to show that (Alg. 7) XClosure gives only correct answers for X^+ w.r.t. Σ, which is known (cf. theorem from Ullman [27, p. 389]). Note that XClosure(Σ, X) inserts A in X^+ only if $A \subseteq X$ or there is some fd $\langle Y, A\rangle \in \Sigma$. □

Lemma 3. *Let Σ be an fd set. An fd $X \to Y$ is in Σ^+ iff $Y \subseteq X^+$, where X^+ is the attribute closure of X w.r.t. Σ.*

Proof. This is from Ullman [27, p. 386]. Let $Y = A_1 \ldots A_n$ and suppose $Y \subseteq X^+$. Then for each A_i, we have $A_i \in X^+$ and, by Def. 5, we must have $\langle X, A_i\rangle \in \Sigma^+$. Then it follows by (R4) union that $X \to Y$ is in Σ^+ as well. Conversely, suppose $\langle X, Y\rangle \in \Sigma^+$. Then, by (R3) decomposition we have $\langle X, A_i\rangle \in \Sigma^+$ for each $A_i \in Y$. □

B.2 Proof of Theorem 1.

"Let Σ be an fd set defined $\Sigma \triangleq$ h-encode($\mathcal{S}$) for some complete structure $\mathcal{S}$. Then Σ is non-redundant but may not be canonical."

Proof. We will show that properties (a-b) of Def. 6 hold for Σ produced by (Alg. 2) h-encode, but property (c) may not hold.

At initialization, the algorithm sets $\Sigma = \varnothing$ and then inserts an fd $\langle X, A\rangle \in \Sigma$ for each $\langle f, x\rangle \in \varphi_t$ scanned, where $x \mapsto A$ and $X \cap A = \varnothing$. At termination, for all fd's in Σ we obviously have $|A| = 1$ then property (a) holds. Also, note that $\varphi_t : S \to Vars(S)$ is, by Def. 11, a bijection.

Now, for property (b) not to hold there must be some fd $\langle X, A\rangle \in \Sigma$ that is redundant and then can be found in the closure of $\Gamma = \Sigma \setminus \langle X, A\rangle$. By Lemma 1, that can be the case only if $A \subseteq X$ or there is $\langle Y, A\rangle \in \Gamma$ for some Y. But from $X \cap A = \varnothing$, we have $A \nsubseteq X$; and from φ_t being a bijection it follows that there can be no such fd in Γ. Thus it must be the case that Σ is non-redundant, i.e., property (b) holds.

Finally, property (c) does not hold if there can be some fd $\langle X, A\rangle \in \Sigma$ with $Y \subset X$ such that $\Gamma = \Sigma \setminus \langle X, A\rangle \cup \langle Y, A\rangle$ has the same closure as Σ. That is, if we may find $\langle Y, A\rangle \in \Sigma^+$. Now, pick structure S whose (3×3) matrix is $A_s = \{1, 0, 0; 1, 1, 0; 1, 1, 1;\}$ as an instance. Alg. 2 encodes it into $\Sigma = \{\phi \to x_1,\ x_1\, v \to x_2,\ x_1\, x_2\, v \to x_3\}$. Let $Y = \{x_1, v\}$, $Z = \{x_2\}$ and $X = YZ$ such that $Y \subset X$. Note that $x_1 v \to x_2 \in \Sigma$ can be written as $\langle Y, Z\rangle \in \Sigma$, and $x_1\, x_2\, v \to x_3 \in \Sigma$ as $\langle YZ, A\rangle \in \Sigma$. Apply R5, R0 to derive $\langle Y, A\rangle \in \Sigma^+$, which is sufficient to show that property (c) may not hold. □

B.3 Proof of Corollary 1.

"*Let Σ be an fd set defined $\Sigma \triangleq$ h-encode($\mathcal{S}$) for some complete structure $\mathcal{S}$. Then Σ can be assumed parsimonious with no loss of generality at expense of time that is polynomial in $|\Sigma| \cdot |U|$.*"

Proof. By Theorem 1, we know that any fd set Σ produced by h-encode is (Def. 6-a) singleton-rhs and (Def. 6-b) non-redundant. Then let $\Sigma' \triangleq$ left-reduce(Σ) (cf. Alg. 8 in §A), which is essentially testing, for all fd's $\langle X, Y\rangle \in \Sigma$ and all attributes $A \in X$, whether $\langle X\backslash A,\ Y\rangle \in \Sigma^+$ Such test is dominated by polynomial-time XClosure (cf. Remark 3). Since it ensures Σ' to be (Def. 6-c) left-reduced, then Σ' must be canonical.

Moreover, in order to verify it is parsimonious, recall that h-encode requires its input structure $\mathcal{S}$ to be complete and then (COA$_t$) processes it into a bijection $\varphi_t : \mathcal{E} \to \mathcal{V}$ from equations to variables. That is, each variable $x \in \mathcal{V}$ is mapped to a relational attribute $x \mapsto B$ and has a unique set of variables $\mathcal{V}' \subset \mathcal{V}$ that is encoded to determine it, i.e., $\Sigma \triangleq$ h-encode$(\mathcal{S})$ is such that $\langle Y, B\rangle \in \Sigma$ for exactly one attribute set Y where $\mathcal{V}' \mapsto Y$. Therefore, Σ' must also be parsimonious. □

B.4 Proof of Proposition 1

"Let $\boldsymbol{R}[U]$ be a relational schema, defined $\boldsymbol{R} \triangleq$ synthesize(Σ) for some canonical fd set Σ on attributes U. Then $\boldsymbol{R}$ preserves Σ, and is in 3NF but may not be in BCNF."

Proof. Note that a relation scheme $R_i[S]$ is rendered into $\boldsymbol{R}$ only if we have an fd $\langle X, Z\rangle \in \Sigma'$ such that $XZ \subseteq S \subseteq U$; and also, for all fd's $\langle Y, W\rangle \in \Sigma'$, there must be some scheme $R_j[T] \in \boldsymbol{R}$ with with $YW \subseteq T \subseteq U$. As fd set Σ is recoverable from Σ' by a finite application of (R3) decomposition (cf. line 2 of Alg. 3), it includes all the fd's projected onto $\boldsymbol{R}$. Since Σ is canonical, by Def. 6 it is a cover for every fd set Γ such that $\Gamma^+ = \Sigma^+$. Then schema $\boldsymbol{R}$ synthesized over Σ' clearly preserves Σ. Now we show that schema $\boldsymbol{R}$ may not be in BCNF. Let Alg. 3 be applied to canonical Σ given in Fig. 4 (left). For synthesized relation $R_2[YW]$, where $Y = \{x_1, x_2, x_3, x_5, \upsilon\}$ and $W = \{x_4\}$, note that both fd's $x_1\, x_2\, x_3\, x_5\, \upsilon \to x_4$ and $x_1\, x_3\, x_4\, \upsilon \to x_5$ hold. But the latter has the form $X \to A$, where $A \not\subseteq X$ but $X \not\to \{x_2\}$ then $X \not\to YW$. That is, by Def. 2, it violates BCNF in R_2 (then in $\boldsymbol{R}$). Note that $A = x_5$ is prime in R_2, thus 3NF is not violated in R_2 by that fd. In fact any $\boldsymbol{R}$ synthesized by Alg. 3 is in 3NF.

To show that $\boldsymbol{R}$ is in 3NF, we first disconsider line 2 of (Alg. 3) synthesize to study the schema as if synthesized over canonical Σ and then extend the proof to Σ'. Let $\langle Y, B\rangle \in \Sigma$ be the fd over which scheme $R_k[YB]$ has been synthesized. We have to prove that R_k is in 3NF. By contradiction, suppose there is some $\langle X, A\rangle \in \Sigma^+$, with $XA \subseteq YB$, that violates 3NF in $R_k[YB]$. That is, by Def. 2, we have $A \not\subseteq X$ but $X \not\to YB$ and A is nonprime. Following Ullman [27, p. 410], we have two cases for analysis. If $A = B$, then $X \to B$ and, since $A \not\subseteq X$ and $XA \subseteq YB$, we have $X \subseteq Y$. But as $X \not\to Y$, it turns out that $X \subset Y$ and $X \to B$, while $Y \to B$ is in supposedly canonical Σ. $\lightning$. Else ($A \neq B$), we have $A \in Y$. Let $Z \subseteq Y$ be a key for YB then, since A is nonprime, $A \notin Z$. That is, we have $Z \subset Y$ and $Z \to B$, while $Y \to B$ is in supposedly canonical Σ. $\lightning$. Therefore $R_k[YB]$ (and $\boldsymbol{R}$ in general) must be in 3NF.

The extension of this proof to Σ' (which is Σ after application of R4) is straightforward. For relation $R_k[YW]$, we may have $W = B_1 B_2 \ldots B_n$, with $Y \to W$ in left-reduced, non-redundant Σ'. Suppose there is some $\langle X, A\rangle \in (\Sigma')^+$, with $XA \subseteq YW$, that violates 3NF for $R_k[YW]$. If $A \subset W$, let $Y' = Y \cup (W \setminus A)$. Then we have $R'_k[Y'A]$ under the same properties analyzed above. Else ($A \notin W$), thus we have $A \in Y$ idem. Either way establishes a contradiction, therefore $R_k[YZ]$ (and $\boldsymbol{R}$ in general) must be in 3NF. □

B.5 Proof of Lemma 2

"Let Σ be a canonical fd set on attributes U, and $\langle X, A\rangle \in \Sigma$ be an fd with $XA \subseteq U$. Then $A^{\looparrowright}$, the attribute folding of A (w.r.t. Σ) exists. Moreover, if Σ is parsimonious then $A^{\looparrowright}$ is unique."

Proof. The existance of $A^{\looparrowright}$ is ensured by the degenerate case where $X = A^{\looparrowright}$, as $X \to A$ is itself in $\Sigma^{\triangleright}$ by an empty application of {R0, R5}. If $X \to A$ is in fact folded w.r.t. Σ, then the folding of A exists. Else, it is not folded yet $X \to A$ is non-trivial because Σ is canonical by assumption. Then, by Def. 15 there must be some $Y \subseteq U$ with $Y \not\supseteq X$ such that $Y \to X$ and $X \not\to Y$. By Def. 14, there is a finite application of rules {R0, R5} over fd's in Σ to derive $Y \to X$. Then by R2 $\sim$ R5 over $X \to A$, we have $Y \to A$. Although there may be many such (intermediate) attribute sets $Y \subset U$ along the causal ordering satisfying the conditions above, we claim there is at least one that is a folding of A. Suppose not. Then, for all such $Y \subset U$, there is some $Y' \subset U$ with $Y' \not\supseteq Y$ such that $Y' \to Y$ and $Y \not\to Y'$, leading to an infinite regress. Nonetheless, in so far as cycles are ruled out by force of Def. 15, then Σ^+ must have an infinite number of fd's. But Σ^+ is finite, viz., bounded by $2^{2|U|}$ (cf. [1, p. 165]). $\lightning$. Therefore the folding of A must exist.

If Σ is assumed parsimonious, then by Def. 13 then we have $\langle X, A\rangle \in \Sigma$ for exactly one attribute set X. Then, as a straightfoward follow-up of the above reasoning that let us infer the folding existance, note that there must be a single chaining $Y^n \overset{\triangleright}{\to} \ldots \overset{\triangleright}{\to} Y^1 \overset{\triangleright}{\to} Y^0 \overset{\triangleright}{\to} X \overset{\triangleright}{\to} A$. Again, as cycles are ruled out by force of Def. 15 and Σ^+ is finite, then the folding of A is unique. □

B.6 Proof of Theorem 2

"Let Σ be a parsimonious fd set on attributes U, and A be an attribute with $\langle X, A\rangle \in \Sigma$ with $XA \subseteq U$. Then AFolding(Σ, A) correctly computes $A^{\looparrowright}$, the attribute folding of A (w.r.t. Σ) in time $O(n^2)$ in $|\Sigma| \cdot |U|$."

Proof. For the proof roadmap, note that AFolding is monotone and terminates precisely when $A^{(i+1)} = A^{(i)}$, where $A^{(i)}$ denotes the attributes in $A^\star$ at step i of the outer loop. The folding $A^{\looparrowright}$ of A at step i is $A^{(i)} \setminus \Lambda^{(i)}$. We shall prove by induction, given attribute A from fd $X \to A$ in parsimonious Σ, that $A^\star \setminus \Lambda$ returned by AFolding(Σ, A) is the unique attribute folding $A^{\looparrowright}$ of A.

(**Base case**). Since Σ is assumed canonical with (then) non-trivial $\langle X, A\rangle \in \Sigma$ for exactly one attribute set X, the algorithm always reaches step $i = 1$, which is our base case. Then X is placed in $A^{(1)}$ and A in $\Lambda^{(1)}$, and we have $A^{(1)} = XA$ and $\Lambda^{(1)} = A$. Therefore, $A^{(1)} \setminus \Lambda^{(1)} = X$, and in fact we have $\langle X, A\rangle \in \Sigma^{\triangleright}$. For it to be specifically in $\Sigma^{\looparrowright} \subset \Sigma^{\triangleright}$, it must be folded w.r.t. set Γ of consumed fd's at this step, viz., $\Delta^{(1)} = \{X \to A\}$. Since $\langle X, A\rangle \in \Sigma^{\triangleright}$ and by Def. 15 is in fact folded w.r.t. $\Delta^{(1)}$, we have $A^{\looparrowright} = X$ at step $i = 1$.

(**Induction**). Let $i = k$, for $k > 1$, and assume that $\langle A^{(k)} \setminus \Lambda^{(k)}, A\rangle \in \Sigma^{\triangleright}$ is in $\Sigma^{\looparrowright}$, with $A^{(k)} \neq \Lambda^{(k)}$. That is, by Lemma 2 we know that $A^{\looparrowright} = A^{(k)} \setminus \Lambda^{(k)}$ is the unique folding of A at step $i = k$. Now, for the inductive step, suppose Y is placed in $A^{(k+1)}$ and B in $\Lambda^{(k+1)}$ because $\langle Y, B\rangle \in \Sigma$ and $B \in A^{(k)}$, with $Y \cap \Lambda^{(k)} = \varnothing$. That is, $A^{(k+1)} = A^{(k)}Y$ and $\Lambda^{(k+1)} = \Lambda^{(k)}B$. We must show that $\langle A^{(k+1)} \setminus \Lambda^{(k+1)}, A\rangle \in \Sigma^{\looparrowright}$.

By the inductive hypothesis, we have $(A^{(k)} \setminus \Lambda^{(k)}) \to A$ in $\Sigma^{\looparrowright} \subset \Sigma^{\triangleright}$. Note that such fd implies $(A^{(k)}B \setminus \Lambda^{(k)}B) \to A$. Since $B \in A^{(k)}$, we actually have $(A^{(k)} \setminus \Lambda^{(k)}B) \to A$. Observe that $(A^{(k)} \setminus \Lambda^{(k)}B) \subseteq (A^{(k)}Y \setminus \Lambda^{(k)}B)$. Then, by R0, we must have $(A^{(k)}Y \setminus \Lambda^{(k)}B) \to (A^{(k)} \setminus \Lambda^{(k)}B)$. Now, by R5, we get $(A^{(k)}Y \setminus \Lambda^{(k)}B) \to A$. Finally, as $A^{(k)}Y = A^{(k+1)}$ and $\Lambda^{(k)}B = \Lambda^{(k+1)}$, we infer $(A^{(k+1)} \setminus \Lambda^{(k+1)}) \to$

A is in $\Sigma^{\triangleright}$. That has been derived by implication from the inductive hypothesis through a finite application of {R0, R5}.

Moreover, with the addition of $Y \to B$ into $\Delta^{(k+1)}$, observe that previous $\langle A^{(k)} \setminus \Lambda^{(k)},\, A\rangle \in \Sigma^{\triangleright}$ is no longer in $\Sigma^{\looparrowright}$, as it is not folded w.r.t. $\Delta^{(k+1)} \supset \langle Y, B\rangle$. In fact, in accordance with Lemma 2, it is replaced by $\langle A^{(k+1)} \setminus \Lambda^{(k+1)},\, A\rangle$ in $\Sigma^{\looparrowright}$, which is folded w.r.t. $\Delta^{(k+1)}$, i.e., $A^{\looparrowright} = A^{(k+1)} \setminus \Lambda^{(k+1)}$.

Finally, as for the time bound, note that in worst case, exactly one fd is consumed into Δ for each step of the outer loop. That is, $|\Sigma| = n$ is decreased stepwise in arithmetic progression such that $n + (n-1) + \ldots + 1 = n\,(n-1)/2$ scans are required overall, thus Alg. 5 is bounded by $O(n^2)$. □

B.7 Proof of Corollary 2

"Let Σ be a parsimonious fd set on attributes U. Then algorithm folding(Σ) correctly computes $\Sigma^{\looparrowright}$, the folding of Σ in time that is $f(n)\,\Theta(n)$ in the size $|\Sigma| \cdot |U|$, where $f(n)$ is the time complexity of (Alg. 5) AFolding."

Proof. By Theorem 2, we know that sub-procedure (Alg. 5) AFolding is correct and terminates. Then (Alg. 4) folding necessarily inserts in Γ (initialized empty) exactly one fd $Z \xrightarrow{\looparrowright} A$ for each fd $X \to A$ in Σ scanned. Thus, at termination we have $|\Gamma| = |\Sigma|$. Again, as AFolding is correct, we know Z is the unique folding of A. Therefore it must be the case that Alg. 4 is correct. Finally, for the time bound, the algorithm iterates in any case over each fd in Σ, and at each such step AFolding takes time that is $f(n)$. Thus folding takes $f(n)\,\Theta(n)$. But we know from Theorem 2 and Remark 3 that $f(n) \in O(n)$, then it takes $O(n^2)$. □

B.8 Proof of Proposition 2

"Let Σ be an fd set, and $\Sigma^{\looparrowright}$ its folding. If Σ is parsimonious then so is $\Sigma^{\looparrowright}$."

Proof. Since Σ is parsimonious, by Lemma 2 we know that, for each fd $\langle X, A\rangle \in \Sigma$, the attribute folding Z of A such that $Z \xrightarrow{\looparrowright} A$ exists and is unique. That is, for no $Y \neq Z$ we have $Y \xrightarrow{\looparrowright} A$. Thus $\Sigma^{\looparrowright} := \mathsf{folding}(\Sigma)$ automatically satisfies Def. 13, as long as we show it is canonical (cf. Def. 6).

First, note that Σ is parsimonious. Then, by Def. 13, it must be canonical and, by Def. 6, it is singleton-rhs, non-redundant and left-reduced. Now, consider by Lemma 2 that AFolding builds a bijection mapping each $\langle X, A\rangle \in \Sigma$ to exactly one $\langle Z, A\rangle \in \Sigma^{\looparrowright}$ such that $Z \xrightarrow{\looparrowright} A$. As Σ is singleton-rhs, it is obvious that $\Sigma^{\looparrowright}$ is as well. Also, the bijection implies $|\Sigma^{\looparrowright}| = |\Sigma|$. Since Σ is non-redundant, then by Lemma 1 so is $\Sigma^{\looparrowright}$.

Besides, suppose by contradiction that $\Sigma^{\looparrowright}$ is not left-reduced. Then there is some fd $Z \to A$ in $\Sigma^{\looparrowright}$ with $W \subset Z$ such that $W \to A$ is in $(\Sigma^{\looparrowright})^{+}$. But as Σ is parsimonious, $X \to A$ is the only fd with A in its rhs and we know $W \not\subseteq X$ by Σ being left-reduced. Then for $W \to A$ to be in $(\Sigma^{\looparrowright})^{+}$, it must be the case that $W \xrightarrow{\triangleright} X$ and then $W \xrightarrow{\triangleright} A$. However, as $Z \xrightarrow{\triangleright} A$ is folded and $W \subset Z$, then $W \xrightarrow{\triangleright} A$ must be folded as well, i.e., the folding of A is not unique. ↯. Therefore $\Sigma^{\looparrowright}$ must be parsimonious. □

B.9 Proof of Theorem 3

"Let $\boldsymbol{R}[U]$ be a relational schema, defined $\boldsymbol{R} \triangleq$ synthesize($\Sigma^{\looparrowright}$) where $\Sigma^{\looparrowright}$ be the folding of parsimonious fd set Σ on attributes U. We claim that $\boldsymbol{R}$ is in BCNF, is minimal-cardinality and preserves $\Sigma^{\looparrowright}$."

Proof. Verification of dependency preservation w.r.t. $\Sigma^{\looparrowright}$ follows just the same logic as in Proposition 1 except that the input fd set is now $\Sigma^{\looparrowright}$, the folding of parsimonious Σ (known to have more specific properties). It must be the case that $\boldsymbol{R}$ preserves $\Sigma^{\looparrowright}$.

Now we concentrate on the BCNF property. We shall use Lemma 4, originally from Osborn [21],[6] to check it in a convenient way. Let $\langle Y, W\rangle \in (\Sigma^{\looparrowright})^{+}$ be the fd over which scheme $R_k[YW]$ was synthesized. We shall prove that R_k[YW] must be in BCNF. By Proposition 2, $\Sigma^{\looparrowright}$ is parsimonious (then canonical), and then by Lemma 4 we only need to check for fd violations in $\Sigma^{\looparrowright}$, not $(\Sigma^{\looparrowright})^{+}$. By contradiction, suppose there is some $\langle X, A\rangle \in \Sigma^{\looparrowright}$, with $XA \subseteq YW$, that violates BCNF in $R_k[YW]$. That is, by Def. 2, we have $A \notin X$ but $X \not\to YW$. Then, as Y is a key for $R_k[YW]$ and $\langle X, A\rangle \in \Sigma^{\looparrowright}$, we have $X \not\subseteq Y$ and then X must (at least) partly intersect with W. So, let $X = ST$ for some $T \neq \varnothing$ with $T \subset W$, and $Y = SZ$ for some $Z \neq \varnothing$ with $Z \cap T = \varnothing$. By assumption, $\langle ST, A\rangle \in \Sigma^{\looparrowright}$. That is, by Def. 15, there can be no $V \not\supseteq ST$ such that $V \to ST$ and $ST \not\to V$. Now, take $V = Y = SZ$. Note that $Z \cap T = \varnothing$ then $SZ \not\supseteq ST$. Since $T \subset W$ and $SZ \to W$, by (R3) decomposition we have $SZ \to T$ and then, by (R0) reflexivity, we get $SZ \to ST$. But $(X \not\to Y)$ $ST \not\to SZ$. That is, $\langle ST, A\rangle \notin \Sigma^{\looparrowright}$. ↯. Thus $R_k[YW]$ (and $\boldsymbol{R}$ in general) must be in BCNF.

For the minimality, note that any two schemes $R_i[YW]$, $R_j[XZ]$ are rendered by synthesize into $\boldsymbol{R}$ iff we have fd's $\langle Y, W\rangle, \langle X, Z\rangle \in (\Sigma^{\looparrowright})'$ and $X \not\leftrightarrow Y$, i.e., it is not the case that both $X \to Y$ and $Y \to X$ hold in $(\Sigma^{\looparrowright})^{+}$. Now, to prove that $\boldsymbol{R}$ is minimal-cardinality, we have to find that merging any such pair of arbitrary schemes shall hinder BCNF in $\boldsymbol{R}$. In fact, take $\boldsymbol{R}' := \boldsymbol{R} \setminus (R_i[YW] \cup R_j[XZ]) \cup R_k[YWXZ]$. As $X \not\leftrightarrow Y$, then neither X nor Y can be a superkey for R_k, i.e., R_k is not in BCNF. □

Lemma 4. *Let $\boldsymbol{R}[U]$ be a relational schema and Σ a canonical fd set on attributes U. Then, $\boldsymbol{R}$ can be verified to be in BCNF by checking for violations w.r.t. Σ only (not w.r.t. Σ^{+}).*

Proof. See Osborn [21]. □

B.10 Proof of Proposition 3

"Let $\boldsymbol{R}[U] \triangleq$ synthesize($\Sigma^{\looparrowright}$) be a relational schema with $|\boldsymbol{R}| \geq 2$, where $\Sigma^{\looparrowright}$ is the folding of a parsimonious fd set $\Sigma \triangleq$ h-encode($\mathcal{S}$) on attributes U. Then $\boldsymbol{R}$ has a lossless join (w.r.t. $\Sigma^{\looparrowright}$) iff, for all $R_i[XZ] \in \boldsymbol{R}[U]$ with key constraint $X \to Z$ for $XZ \subseteq U$, we have $X \to U$ or there is some scheme $R_j[YW] \in \boldsymbol{R}[U]$ with key constraint $Y \to W$ for $YW \subseteq U$ such that $X \subset Y$."

[6] Cf. also Ullman [27, p. 403].

PROOF. We use Lemma 5 from Ullman [27, p. 397], which gives a convenient necessary and sufficient condition for the lossless join property. By Lemma 5, any pair $R_i[XZ], R_j[YW] \in \boldsymbol{R}[U]$ with $XZYW \subseteq U$ have a lossless join w.r.t. $\pi_{XZYW}(\Sigma^{\leftrightarrow})$, iff $(X \cap Y) \rightarrow (XZ \setminus YW)$ or $(X \cap Y) \rightarrow (YW \setminus XZ)$ hold in $\pi_{XZYW}((\Sigma^{\leftrightarrow})^+)$.

Now, let $R_i[XZ], R_j[YW] \in \boldsymbol{R}$ be arbitrary schemes. By Theorem 3, $\boldsymbol{R}$ is in BCNF. Then, if $R_i[XZ]$ and $R_j[YW]$ are two different schemes in $\boldsymbol{R}$, we must have $XZ \setminus W = XZ$ and $YW \setminus Z = YW$. That is, we can write the condition imposed by Lemma 5 as $(X \cap Y) \rightarrow (XZ \setminus Y)$ or $(X \cap Y) \rightarrow (YW \setminus X)$.

Suppose $X \subset Y$. That is, we have $X \cap Y = X$ and $X \rightarrow XZ$ then obviously $X \rightarrow XZ \setminus Y$. That is, $R_i[XZ]$ and $R_j[YW]$ have a lossless join. Now, suppose rather that $X \rightarrow U$. Then $X \rightarrow Y$ but $Y \not\rightarrow X$ and $Y \neq X$, otherwise $R_i[XZ]$ and $R_j[YW]$ would have been merged by synthesize. Moreover, by (R2) transitivity we have $X \rightarrow W$. But for some $B_i \in W$ we must have had $Y \xrightarrow{\leftrightarrow} B_i$ in $\Sigma^{\leftrightarrow}$. Since we do have $X \rightarrow Y$ and $Y \not\rightarrow X$, then by Def. 15 it must be the case that $X \supset Y$. Therefore $X \cap Y = Y$ and $Y \rightarrow YW$, then obviously $Y \rightarrow YW \setminus X$. Since R_i and R_j are picked arbitrarily and the natural join operator is associative, $\bowtie_{R_\ell \in \boldsymbol{R}} \pi_{R_\ell}(r)$ must be lossless.

For the converse, suppose rather that $X \not\rightarrow U$ and for all such $R_j[YW]$ we have $X \not\subset Y$. Since $X \neq Y$, we actually have $X \not\subseteq Y$. Now, let $X = ST$ and $Y = SV$ for some $T \neq V$ and $T \neq \varnothing$. That is, $X \cap Y = S$, and (i) $R_i \setminus R_j = STZ \setminus SVW = TZ$ but we have $S \not\rightarrow Z$ (otherwise T would be extraneous in key X); and (ii) $R_j \setminus R_i = VW \setminus TZ = VW$ but we have $S \not\rightarrow V$ likewise. Then there can be no $R_j[YW]$ which $R_i[XZ]$ has a lossless join with, and thus $\boldsymbol{R}[U]$ cannot have a lossless join. □

LEMMA 5. *Let Σ be a set of fd's on attributes U, and $R_i[S], R_j[T] \in \boldsymbol{R}[U]$ be relation schemes with $ST \subseteq U$; and let $\pi_{ST}(\Sigma)$ be the projection of Σ onto ST. Then $R_i[S]$ and $R_j[T]$ have a lossless join w.r.t. $\pi_{ST}(\Sigma)$ iff $(S \cap T) \rightarrow (S \setminus T)$ or $(S \cap T) \rightarrow (T \setminus S)$ hold in $\pi_{ST}(\Sigma^+)$.*

PROOF. See Ullman [27, p. 397]. □

B.11 Proof of Proposition 4

"Let $H^{\ell}_k[XZ] \in \boldsymbol{H}_k[U]$ be an exogenous relation with (violated) key constraint $\langle X, Z\rangle \in \Sigma^{\leftrightarrow}_k$ for $XZ \subseteq U$. Then,

- *(a) for any pair $Y^i_k[\,V_iD_i \,|\, XA_i\,]$, $Y^j_k[\,V_jD_j \,|\, XA_j\,]$ of u-factor projections of H^{ℓ}_k, they are independent.*
- *(b) the join $\bowtie^m_{i=1} Y^i_k[\,V_iD_i \,|\, XA_i\,]$ of all u-factor projections of H^{ℓ}_k is lossless w.r.t. $\pi_{XA_1\,A_2\,\ldots\,A_m}(\Sigma^{\leftrightarrow}_k)$."*

PROOF. We prove the claims in separate as follows.

(a) Suppose not. Then we have either $A_i \rightarrow A_j$ or $A_j \rightarrow A_i$. Let G_i be such that $A_i \in G_i \subseteq Z$. Then, since $Y^i_k[\,V_iD_i \,|\, XA_i\,]$ is a u-factor projection of $H^{\ell}_k[XZ]$, by Def. 18 there can be no $C \in Z \setminus G_i$ with $A_i \rightarrow C$ or $C \rightarrow A_i$. Now, take G_j such that $A_j \in G_j \subseteq Z$. Since $G_i, G_j \subseteq Z$ are learned to be maximal groups that have to be disjoint in Z (cf. Problem 4), we must have such $A_j \in Z \setminus G_i$ with $A_i \rightarrow A_j$ or $A_j \rightarrow A_i$. ↯.

(b) First, note that the lossless join property is considered w.r.t. $\pi_{XA_1\,A_2\,\ldots\,A_p}(\Sigma^{\leftrightarrow}_k)$. By Lemma 5, we know that any pair $Y^i_k[\,V_iD_i \,|\, XA_i\,]$, $Y^j_k[\,V_jD_j \,|\, XA_j\,]$ of u-factor projections will have a lossless join w.r.t. $\pi_{XA_i\,A_j}(\Sigma^{\leftrightarrow}_k)$ iff $(XA_i \cap XA_j) \rightarrow (XA_i \setminus XA_j)$ or $(XA_i \cap XA_j) \rightarrow (XA_j \setminus XA_i)$ hold in $\pi_{XA_i\,A_j}((\Sigma^{\leftrightarrow}_k)^+)$. By Def. 18, we must have $XA_i \cap XA_j = X$, and $XA_i \setminus XA_j = A_i$. In fact $X \rightarrow A_i$ is in $\pi_{XA_1\,A_2\,\ldots\,A_p}(\Sigma^{\leftrightarrow}_k)$, thus Y^i_k and Y^j_k have a lossless join w.r.t. such fd set projection. Then, from the associativeness of the join, all the u-factor projections taken together must have a lossless join. □

B.12 Proof of Theorem 4

"Let $H^q_k[Z_qV] \in \boldsymbol{H}_k$ be an endogenous relation with (violated) key constraint $\langle Z_q, V\rangle \in \Sigma^{\leftrightarrow}_k$, and $Y^r_k[\overline{V_iD_i} \,|\, Z_qT]$ be a predictive projection of H^q_k w.r.t. $\Gamma^{\leftrightarrow}_k$ defined by formula (6) with $V \supseteq T$. We claim that Y^r_k correctly captures all the uncertainty present in H^q_k w.r.t. $\Gamma^{\leftrightarrow}_k$."

PROOF. For the proof roadmap, consider that violated key constraint $\langle Z_q, V\rangle \in \Sigma^{\leftrightarrow}_k$ is 'trivially' repaired '4C' by using special attribute 'trial id' tid provisionally, viz., take $Z'_q = Z_q \cup \{\mathsf{tid}\}$ and then $Z'_q \rightarrow V$ holds in endogenous relation $H^q_k[Z'_qV]$. Exogenous relations of form $H^{\ell}_k[X_\ell\, L]$ are trivially repaired likewise, viz., $X'_\ell := X_\ell \cup \{\mathsf{tid}\}$ to get clean $H^{\ell}_k[X'_\ell\, L]$, so that we have (i) a clean loading of trial data, and (ii) a (correct) natural join of tuples from exogenous and endogenous relations. In such setting '4C,' tid works for every endogenous tuple as a surrogate to the exogenous tuple(s) that are in its causal chain. In fact, what u-propagation does in synthesis '4U' is to replace tid by some set $\overline{V_iD_i}$ of pairs of condition columns such that, for $Z''_q = \overline{V_iD_i}Z_q$, we know that $Z''_q \rightarrow T$ holds in U-relation $Y^r_k[Z''_q\,T]$, a predictive projection of $H^q_k[Z'_qV]$ w.r.t. $\Gamma^{\leftrightarrow}_k$. By Def. 19, query formula (6) we outline below accordingly defines exactly such replacement, and what we shall prove by construction is that it does it correctly w.r.t. $\Gamma^{\leftrightarrow}_k$.

$$Y^r_k \;:=\; \pi_{Z_qT}(\,\sigma_{\upsilon=k}(Y_0) \bowtie \pi_{X'}(J) \bowtie \pi_{Z'_qT}(H^q_k)\,)$$

where $X' = \mathsf{tid} \cup X$ and $Z'_q = \mathsf{tid} \cup Z_q$.

We start by considering (as the join is associative) sub-query $\sigma_{\upsilon=k}(Y_0) \bowtie \pi_{Z'_qT}(H^q_k)$. By the rewriting rule for the product on U-relations (cf. §2.1), it is rewritten to be applied over $\sigma_{\upsilon=k}(Y_0[V_0D_0 \,|\, \phi\,\upsilon])$ and $\pi_{Z'_qT}(H^q_k)$, where $\upsilon \in Z'_q$. Moreover, since Σ_k is asumed to have empirical grounding, we must have $\phi \in Z'_q$ as well. The partial result set then is $Q_1[V_0D_0 \,|\, Z'_qT]$.

Now we consider $\pi_{X'}(J)$, where J is a join sub-query defined according to Def. 19. Before we proceed to examine J, though, recall that $Y^r_k[\overline{V_iD_i} \,|\, Z_qT]$ is by assumption a predictive projection of $H^q_k[Z_qV]$. Then there must be $S \supseteq Z_q \setminus \{\phi\}$, where $\langle S, T\rangle \in (\Gamma^{\leftrightarrow}_k)^+$ is the key constraint of sketched scheme $R[ST]$ with $V \supseteq T$. Note that S contains the u-factors for T, unfolded out of their compact representation by $\phi \in Z_q$ such that $\langle Z_q, V\rangle \in (\Sigma^{\leftrightarrow}_k)^+$. Then, by Def. 19, and taking advantage again of the associativeness of the join,

we have:

$$J = (\bowtie_{H_k^\ell \in \mathsf{M}} H_k^\ell) \bowtie (\bowtie_{Y_k^j \in \mathsf{M}(H_k^\ell)} (Y_k^j)).$$

where M is the mapping, from each exogenous relation of form $H_k^\ell[X'_\ell\, L]$ with $(L \cap S) \neq \varnothing$, to the set of all their u-factor projections $Y_k^j[V_j D_j | X_\ell\, A_j]$ such that we have $A_j \in (L \cap S)$ and $X = \bigcup_{H_k^\ell \in \mathsf{M}} X_\ell$.

Note that, by Proposition 4, the join $\bowtie_{Y_k^j \in \mathsf{M}(H_k^\ell)} (Y_k^j)$ is lossless. Thus, by the rewriting on U-relations another partial result set is $\pi_{X'}(J) = Q_2[\overline{V_j D_j}|X']$, where $\overline{V_j D_j}$ is the set of pairs of condition columns containing the random variables associated with all the empirical u-factors in S and that we want to propagate into $Y_k^r[\overline{V_i D_i}\,|Z_q T]$, i.e., $\overline{V_i D_i} = V_0 D_0 \cup \overline{V_j D_j}$. In fact, back to the formula (6),

$$Y_k^r := \pi_{Z_q T}(Q_1[V_0 D_0 \,|\, Z'_q T] \bowtie Q_2[\overline{V_j D_j}|X'].$$

We expect this join to be lossless. By Proposition 3 that will be the case iff $X' \subset Z'_q$, i.e., $X \subset Z_q$, which we show next. In fact, recall that we have $S \to T$ in $(\Gamma^{\looparrowright})^+$ and $Z_q \to T$ is an v-fd in $(\Sigma^{\looparrowright})^+$. Let $B \in T$, then $S \to B$ is in $\Gamma^{\looparrowright}$ and $Z_q \to B$ is in $\Sigma^{\looparrowright}$. That is, $Z_q \xrightarrow{\looparrowright} B$ is folded. For each key constraint $X_\ell \to A_j$ such that X_ℓ was built into X, note that $A_j \in S$ and $X_\ell \to A_j$ is a ϕ-fd in $\Sigma^{\looparrowright}$, i.e., $X_\ell \xrightarrow{\looparrowright} A_j$ is folded as well. That is, X_ℓ must have replaced A_j into Z_q. That is, we must have $X \subset Z_q$.

Therefore the join $Q_1[V_0 D_0 \,|\, Z'_q T] \bowtie Q_2[\overline{V_j D_j}|X']$ is lossless w.r.t. $\Gamma_k^{\looparrowright}$ and we have the result set materialized into $Y_k^r[V_0 D_0 \overline{V_j D_j}\ Z_q T]$, which has exactly the u-factors indicated by $\langle S, T\rangle \in (\Gamma_k^{\looparrowright})^+$. That is, u-propagation into Y_k^r according to formula (6) must be correct. □

C. DETAILED COMMENTS

C.1 Comments on Conjecture 1

As previsouly suggested in §3, our framework brings forth a translation of SEM's concepts into the language of fd's. In general, that seems to be an interesting result, as an expression of properties of deterministic hypothesis in relational theory. Along these lines, say, Proposition 5 below relates to Def. 17.

PROPOSITION 5. *Let Σ be an fd set over attributes U, defined $\Sigma :=$ h-encode($\mathcal{S}$) for some complete structure $\mathcal{S}(\mathcal{E}, \mathcal{V})$, and $x \mapsto A$ for some $x \in \mathcal{V}$ and $A \in U$. Then A is exogenous (endogenous) iff x is exogenous (endogenous).*

PROOF. This shall be straightforward from Def. 12 and Def. 17, by observing (Alg. 2) h-encode. □

Conjecture 1 works in the converse direction towards more complete results in terms of the equivalence between SEM's and our design-theoretic framework.

"The lossless join property is reducible to the structure $\mathcal{S}$ given as input to the pipeline."

We leave to future work to show that the condition for the lossless property (cf. Proposition 3) in fact translates into the structure $\mathcal{S}$ given as input to the pipeline. Note that we have technical means to decide/test for the lossless join efficiently: Ullman gives a polynomial-time algorithm for that [27, p. 406], and Proposition 3 gives an even more specific condition (as the schema is known to be in BCNF) that can be fruitfully exploited in that sense. Thus, we shall be able to decide whether a given structure $\mathcal{S}$ is "fully interconnected" by processing its folding and testing its rendered schema for the lossless join.

Moreover, consider fd set Γ from Example 4. Arguably, it may not make sense in practice to keep those two pairs of variables in the same hypothesis/model — the modeler user would likely strive for each hypothesis to form a single, cohesive piece of scientific inquiry. That is, we expect in practice for the majority of hypotheses given/extracted that they shall have a lossless join. In fact, we find some evidence for that from our initial experiments in a real-world use case (cf. §6.2).